

Information Maximization Fails to Maximize Expected Utility in a Simple Foraging Model

Edward K. Agarwala^{1,2}, Hillel J. Chiel^{3,4,5}, and Peter J. Thomas^{1,3,6,7}.

¹⁻⁶Case Western Reserve University, Departments of

¹Mathematics, ²Operations Research, ³Biology,

⁴Neurosciences, ⁵Biomedical Engineering, and ⁶Cognitive Science.

⁷Oberlin College Department of Neuroscience.

October 25, 2010

Abstract

Information theory has successfully explained the organization of many biological phenomena, from the physiology of sensory receptive fields to the variability of certain DNA sequence ensembles. Some scholars have proposed that information should provide the *central* explanatory principle in biology, in the sense that any behavioral strategy that is optimal for an organism's survival must necessarily involve efficient information processing. Here we challenge this view by providing a counterexample. We present an analytically tractable model for a particular instance of a perception-action loop: a creature searching for a wandering food source confined to a one-dimensional ring world. The model incorporates the statistical structure of the creature's world, the effects of the creature's actions on that structure, and the creature's strategic decision process. The underlying model takes the form of a Markov process on an infinite dimensional state space. To analyze it we construct an exact coarse graining that reduces the model to a Markov process on a finite number of "information states". This mathematical technique allows us to make quantitative comparisons between the performance of an information-theoretically optimal strategy with other candidate search strategies on a food gathering task. We find that

1. Information optimal search does not necessarily optimize utility (expected food gain).
2. The rank ordering of search strategies by information performance does not predict their ordering by expected food obtained.
3. The relative advantage of different strategies depends on the statistical structure of the environment, in particular the variability of motion of the source.

We conclude that *there is no simple relationship between information and utility*. Behavioral optimality does not imply information efficiency, nor is there a simple tradeoff between the two objectives of gaining information about a food source *versus* obtaining the food itself.

Keywords: Infotaxis, Ideal Observer, Markov Process, Lumping, Coarse Graining, Search Strategy, Perception-Action Loop.

1 Introduction

1.1 Information optimality in sensory systems

It has long been recognized that information efficiency is an important explanatory principle underlying the function of biological sensory systems [4, 5]. Linsker [20] proposed that an information maximization principle might govern the structure of layered self-adaptive neural networks such as those found in the visual system. It was subsequently shown that efficient encoding of the visual and auditory environment could account for receptive field physiology observed in cats and macaques [35, 39]. Analysis of information optimality in sensory systems typically proceeds independently of considerations of the relative biological significance of different stimuli in the environment. Simoncelli and Olshausen [34] noted that the principle of efficient coding succeeds in part because it does *not* take into account complicating factors such as the accuracy with which signals are represented, or the costs to the organism of making mistakes.

Several authors have extended the framework of information theory to take into account the biological context of behaving organisms. Shraiman and colleagues, for instance, coined the term *infotaxis* [40] to describe search strategies that locally maximize the expected rate of information gain by an organism. In the context of a model organism pursuing the source of a diffuse pheromone plume, these authors showed that infotaxis outperformed chemotaxis (locally maximizing the expected rate of increase in signal concentration) with respect to mean latency of capture.

Information theoretic arguments have succeeded in molecular biology as well [25]. Schneider, for instance, demonstrated that the statistical entropy of nucleic acid sequences encoding location information read out by proteins docking at specific DNA sites can be predicted by the amount of information required to specify the docking sites among all possible binding sites, *i.e.* the mean entropy decrease from the undocked to the docked state [28, 29]. Moreover the uncertainty decrease in such “molecular machines” can be related *quantitatively* to the free energy of protein–DNA binding [30].

These results underscore the importance of information theory for understanding sensory, behavioral and genetic systems. Some authors go so far as to suggest that information theory could serve as an overriding explanatory principle throughout biology. For example, Adami writes

The discovery of the genetic code cemented the fact that information is the *central* pillar in any attempt to understand life, and the dynamics of information storage and acquisition that come with it. [[1], pg. 59, original emphasis].

In a similar spirit Polani has proposed parsimonious information processing as a universality principle in biology, arguing that “if organisms would develop a suboptimal information-processing strategy, this would waste metabolic energy. Such a disadvantage would then be selected against by evolution.” [[26], p.3]. For the sake of argument we will state what we call the *strong hypothesis*:

Hypothesis 1.1. *Behaviors that are optimal for an organism’s survival are necessarily information optimal.*

Evaluating such a hypothesis requires precise definitions of behavioral and information optimality. On its face, the hypothesis is plausible, inasmuch as it would appear that any successful survival strategy that led to suboptimal information processing should be capable of improvement through improved efficiency of information processing. However, it is possible that complicating factors such as metabolic constraints or the costs of acquiring and processing information may play a role on a par with that of information maximization *per se*. We undertook a quantitative analysis of the relationship between information optimality and utility by constructing an analytically tractable model of a creature using sensory information to search for a source of food, taking care to allow for precise definitions of behavioral and information optimality. As we shall show in the sequel, we find that *even in the absence of metabolic and computational constraints*, information-optimal behavior and utility-optimal behavior bear no simple relation to one another.

1.2 Perception-Action Loop

Studies of the interplay of sensory processing and behavior are complicated by the structure of the problem, which necessarily embodies a closed “perception-action loop”. Barlow’s redundancy reduction hypothesis [5] led to models of sensory processing within the framework now known as acyclic graphical models [16]. In such models the pattern of statistical dependencies form an unambiguous chain or tree from causes (*e.g.* environmental signals) to effects (*e.g.* activity in a layer of sensory neurons). The analysis of information efficiency in this context translates into statements about chains of self consistent conditional probabilities relating the distributions of random variables defined on the nodes of the acyclic graph; one obtains the optimal architecture to accomplish a processing goal (*e.g.* output decorrelation) given the statistical structure of the sensory world. But although modeling sensory perception may be approached along such lines, allowing for a creature’s behavior requires a fundamentally different probabilistic framework.

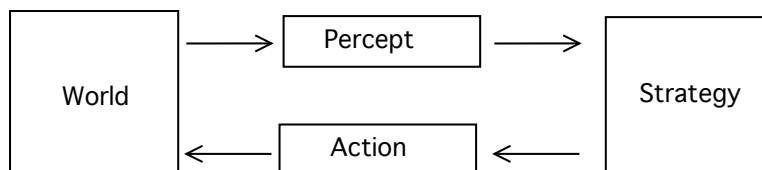


Figure 1.1: The perception-action loop. The statistical structure of the world determines the percept; the creature determines an action by virtue of a strategy (taking into account the percept); the action changes the statistical structure of the world, relative to the creature.

Consider a scenario in which a creature observes the environment in order to accomplish some survival related goal, such as locating and exploiting a food source. Having observed the available stimuli a creature may take action that in turn alters the statistical structure of the world, as seen from the creature’s perspective. In the simplest case, the movement of the creature towards or away from the source changes the statistics of the target related stimuli ¹. Under these circumstances the conditional probabilities describing the stimulus (the world), the received signal (a percept),

¹An analogous situation arises in saccadic visual search, because the spatial resolution of the retina varies significantly from the foveal to the peripheral region. Making a saccade both gathers additional information about possible target locations and changes the statistical distribution of the inputs relative to the (unknown) target location [23]. In both that situation and in the model considered here, the problem becomes tractable when placed in an ideal observer framework [11].

the creature’s decision about how to proceed (a strategy) and the effects of the creature’s behavior (an action) form a closed loop (Figure 1.1). Perception-action loops appear in models of behavior ranging from simplified computational models of evolving agents [17] to conceptual models of cognitive processing in the mammalian prefrontal cortex² [10], as well as systems in which the brain, the body and the environment coevolve [7]. The importance of feedback loops reciprocally connecting a creature’s “input” and “output” has been recognized since the earliest attempts to apply information theory to biology (*cf* [3], Chapter 3).

As is well known, statistical problems such as inference, estimation and sampling, which are well understood for probabilistic models on acyclic graphs, become essentially intractable for graphical models containing cyclic dependencies [41, 42]. The difficulty of analyzing probabilistic systems on graphs with cycles has been identified as a significant barrier to the progress of theory in both neuroscience and computational intelligence [6]. We overcome this obstacle for a relatively tractable model system by incorporating the entire system of interest into a larger Markov process for which the full transition matrix and steady state distribution can be obtained analytically. To this end, we consider a simple model of an organism searching for a target, introduced below (Section 1.3). We begin with a creature that can approximately estimate the location of a food source by measuring the (noisy) local food concentration at a sequence of locations. The source itself moves unpredictably, following a random walk. The combination of the creature and the world amounts to a Markov random walk in a large state space. We then obtain a second, even more tractable model from the first through a heuristic limiting process, by reducing the noisiness of the food concentration measurements. In the low noise limit we are able to reduce the entire system to a Markov process on a small number of states. For the reduced system we then calculate the quantities of interest exactly: the mean amount of food gathered and the average uncertainty about the source location, and how these quantities depend on the creature’s choice of search strategy. We also vary a key parameter, the variability of the motion of the food source itself, and study its effect on the creature’s performance under different strategies.

An important feature of our system is the inclusion of an internal state in the model organism. Whether they involve memory or metabolism, internal states play an important role in many models of biological decision making. Life history theory, to cite one example, incorporated internal states as a form of phenotypic plasticity, thereby increasing its explanatory power. Inclusion of information about the metabolic and environmental resources available to an organism improved understanding of decisions such as reproductive timing [22]. A similar distinction underlies the difference between chain reflex models of locomotion [15, 33], and robust interplay of a central pattern generator (internal state) guided by external information [19, 21]. An analogous dichotomy appears in engineering and control theory, namely the distinction between open-loop control (throwing a baseball as precisely as possible at a target) and closed-loop control (flying an airplane with sensory feedback providing ongoing course corrections). Notably, Touchette and Lloyd [38] analyzed both open- and closed-loop control in a general information theoretic framework. In a control theoretic setting the requirement of controlling the state of a system while also obtaining information about unknown parameters describing it leads to the challenging “dual-control” problem [8, 9, 13].

In the model we present here all four parts of a perception-action loop are represented in nontrivial terms, yet the entire system remains tractable enough to yield a complete analytic treatment. This approach allows us to obtain quantitative results on the success of different search strategies. We

²Compare Figure 1 of [10] with Figure 1.1 of the present manuscript.

consider three movement algorithms, defined below (Section 1.4). Each corresponds to a particular search strategy: one designed to minimize the creature’s *uncertainty about the location* of the food source (the “Information Theory Creature”, ITC), one designed to maximize the *probability of colocation* of the creature with the food source (the “Maximum Likelihood Creature”, MLC) and a third strategy combining elements of the first two (the “Hybrid” or “modified Maximum Likelihood Creature”, mMLC). We studied the performance of each algorithm as a function of a parameter controlling the unpredictability of the source movement. As expected, the ITC had more information about the source location on average than the other two, for all values of the source mobility parameter (Figure 2.9). Surprisingly, however, no one search strategy consistently dominated the others in terms of utility. Each strategy earned a larger expected food yield than the other two for some range of mobility parameters (Figure 2.11). The ITC performed the best when the movement of the food source was the most random, yet the mMLC that used a mixture of this information-optimal strategy and the maximum likelihood based strategy performed the best when the source was the most predictable. The MLC dominated the others for intermediate source mobility values.

In short we show that for this tractable model system there is *no obvious simple relationship* between the information theoretic quality of a search strategy and its performance in terms of expected utility. It therefore serves as a counterexample against the strong hypothesis 1.1. Our approach – embedding the creature and its world into a single Markov process – can be generalized in a variety of ways to incorporate additional biological details. We consider several possible extensions in Section 3.

1.3 Two Models of Search with Perfect Information

Appendix A provides a summary of notation and abbreviations used throughout.

1.3.1 Features common to both models

In both models, the source (or injector) and the creature inhabit a discrete ring world with N sites, identified with \mathbb{Z}_N . The source location $S_t \in \{0, \dots, N-1\}$ is a random variable that is a function of discrete time $t \in \{0, 1, 2, \dots\}$. Distance between locations on the ring is the minimum number of steps clockwise (CW) or counter-clockwise (CCW) between the locations. Formally, the distance between locations l and m is equal to $\min\{l - m \bmod N, m - l \bmod N\}$ (see Figure 1.2).³

The source performs a random walk with uniformly distributed initial position S_0 and independent increments. (Without loss of generality, we set the creature’s initial position, C_0 , to zero.) The source’s walk is governed by a probability mass function (PMF) P , defined such that the probability of the source moving m steps CW is $P(m)$. We denote the set of PMFs on \mathbb{Z}_N , also known as the probability simplex on \mathbb{Z}_N , as $\mathcal{P}(\mathbb{Z}_N)$. The probability mass function for the source location at time $t+1$, S_{t+1} is obtained by convolving P with the PMF for the current location S_t . For simplicity we consider movements to the right or left of the current position equally likely, and thus source

³Addition and subtraction of positions on the ring are interpreted mod N throughout. By convention, we will enumerate loci on the circle clockwise. For simplicity we take N to be even.

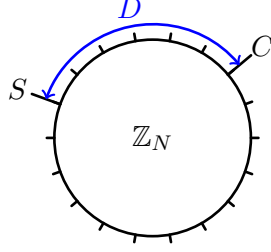


Figure 1.2: Geometry of the ring world, \mathbb{Z}_N , occupied by the creature and the food source. C is the creature's location. S is the source's location. D is the distance between C and S , taken to be $D = \min\{S - C \bmod N, C - S \bmod N\}$.

movement distributions take the form

$$P_x(m) = \begin{cases} x, & m = 1 \\ 1 - 2x, & m = 0 \\ x, & m = -1. \end{cases} \quad (1.1)$$

For technical reasons we will restrict the source mobility parameter x to the range $1/4 < x < 1/3$ (see Figure 2.2).

The world evolves following discrete time dynamics. At each time t , the source establishes a distribution of food molecules around the ring with a diffusion and decay process producing independent Poisson distributed molecule counts at each site. The means are determined by an equilibrium condition balancing an injection rate γ , a decay rate α and a transition rate y between adjacent nodes. Let λ_m , $m \in \{0, \dots, N-1\}$ be the mean number of molecules m steps CW of the source, and define the vector of means $\Lambda = (\lambda_0, \dots, \lambda_{N-1})$. As calculated in Appendix B, the vector Λ is obtained (numerically) by solving a linear system of equations. Figure 1.3 illustrates a typical food distribution; in this case $\alpha = .1$ and $y = 1/3$.

The means of the food distribution follow a strictly decreasing function of distance. Therefore the average amount of food is always higher nearer the injector and the expected amount food m steps CW or CCW of the injector is the same. Additionally, the vector of means Λ is directly proportional to the injection rate γ . Results for other food distributions are discussed in Section 2.

A creature decides how to move based on its observations. The random variable C_t represents the creature's location at time t , and the random variable Z_t represents the creature's observation at the location C_t at time t . We consider two types of observations depending on the model, namely, molecule counts (in Model I) and distance to the source (in Model II). Section 1.3.2 below describes these two scenarios in more detail. In both cases, the current observation Z_t , given the creature's and source's locations at time t , is independent of past observations or locations.

We assume the creature is able to maintain a history of past locations and observations, for instance as a form of working memory. In Appendix D we introduce a formal definition of a history, parallel to the concept of a sigma algebra in the theory of stochastic processes, in order to obtain the results presented in Section 2.1. For ease of presentation here, however, it suffices to represent a history informally as a growing list of observations $Z_t = z_t$ and creature locations $C_t = c_t$. Therefore we

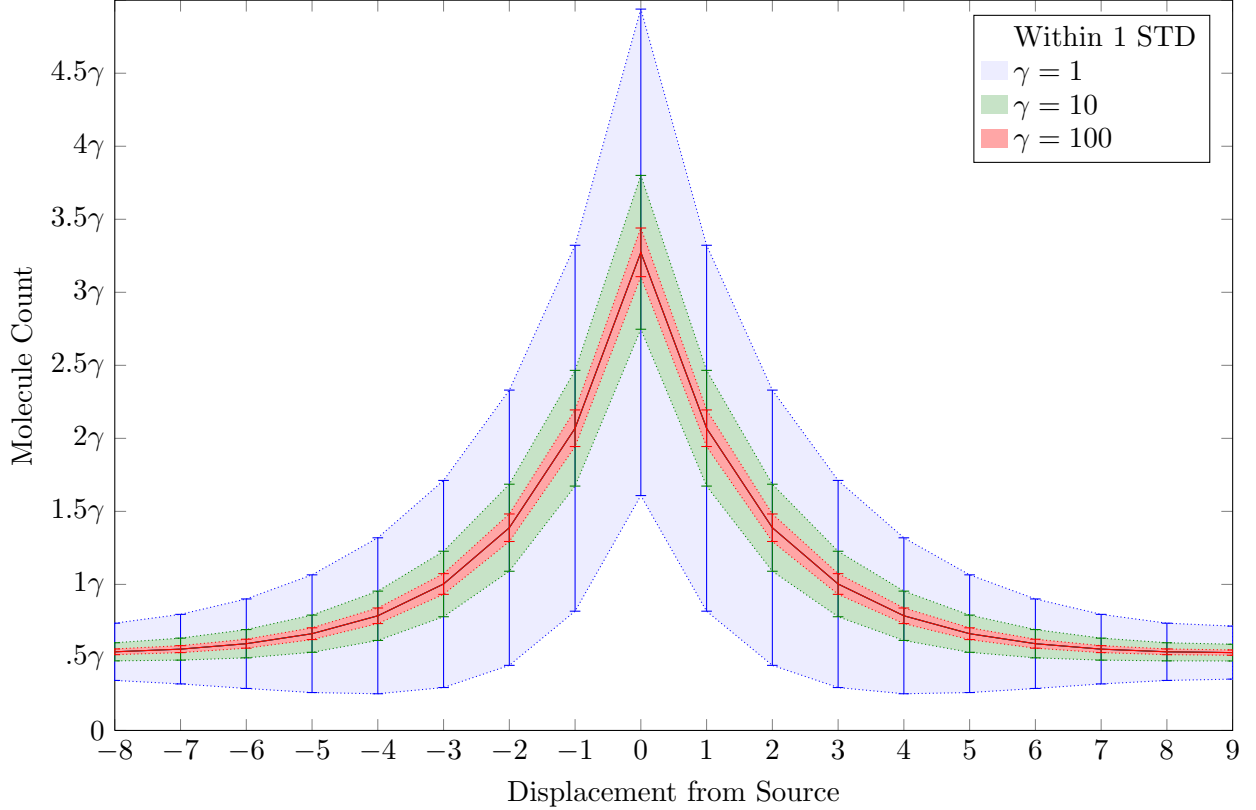


Figure 1.3: Typical food distribution. The mean number of food particles at each location relative to the injector as a function of the injection rate γ , for parameter values $\alpha = .1$, $y = 1/3$. The error bars represent one standard deviation about the mean. Note the vertical axis is scaled according to the injection rate γ . As γ increases the mean number of food particles a given distance from the source increases linearly in γ , while the standard deviation increases as $\gamma^{1/2}$, so the ratio of the standard deviation to the mean decreases. Hence the *shape* of the plot is independent of the injection rate, but the *number of standard deviations* between adjacent means increases without bound as γ goes to infinity. Increasing α or decreasing y gradually narrows the distribution. Subsequent results do not depend strongly on the choice of these parameters.

define a creature’s *history at time t* to be the list of observations and creature locations up to and including time t . A history is then $\{\{Z_i = z_i\}, \{C_i = c_i\}\}_{i=0}^t$.⁴ We define a creature’s *strategy* to be a (possibly random) mapping from histories to a choice for the creature’s next location. The strategies under consideration are discussed in Section 1.4.

The model creature will be endowed with perfect information about several aspects of its environment. In order to focus on the core issue of the (in)equivalence of information maximization and utility maximization, we will simplify the problem by assuming the creature has complete knowledge about every aspect of the world *except* the injector’s actual location. Therefore we assume the creature knows the following:

1. The creature knows that there are exactly N different locations on a 1-dimensional, ring world.
2. The creature knows that at time $t = 0$ the injector starts at a uniformly random location.
3. The creature knows the history of its own locations and observations.
4. The creature knows the injector’s movement algorithm, P_x , and the value of x .
5. The creature knows that the molecules are distributed as Poisson random variables with known means Λ centered on the (unknown) source location.

In addition, we assume the creature’s movement is unconstrained, *i.e.* it can move to any position on the ring at the next iteration, without limitation. This assumption greatly simplifies the formulation and analysis of strategies, as well shall see. It is not entirely unreasonable, however, as in some instances animals may travel substantial distances between foraging or search locations, a behavior known as *saltatory search* or *intermittent locomotion* [18, 24, 27]. Similarly, saltatory movements unconstrained by distance are observed in the pattern of saccades during visual search [23].

Assumption 3 above would appear to require that the foraging creature retain an infinitely capacious memory of all of its prior locations and observations. For the strategies under consideration, however, we will show in Section 2.1 that if the creature has the ability to maintain an internal state representing a single probability mass function on \mathbb{Z}_N , this knowledge is equivalent to the creature knowing its full history. Hence only a finite dimensional “memory” is required in order to exactly encapsulate the creature’s entire history in a natural way.

1.3.2 Discrete Molecule and Distance-Certain Models

In the Discrete Molecule model (Model I), the creature observes the number of molecules at its location at time t . The Poisson distribution of food at each location relative to the injector is conditionally independent of past observations, given the current location of the source. Therefore given the source’s and creature’s locations, the creature’s observation is independent of past observations. The independence of successive observations and the creature’s assumed knowledge of the

⁴Unless otherwise stated, all references to histories are made with respect to time t .

world allows the creature to determine the probability mass function describing the injector’s next location accurately (see Section 2.1).

As the rate of food injection γ increases, the mean number of molecule counts at each location grows linearly in γ while the standard deviation about the mean grows as $\sqrt{\gamma}$. Given an observation of molecule counts at time t , the creature can estimate the distance to the source because it knows the shape of the distribution of food relative to the source. The food distribution is symmetric about the source; once released, food is presumed to move *via* diffusion without drift (see Appendix B). The creature can therefore only estimate the absolute distance, not the direction to the source, from the observed molecule count. However, it can be argued [2], using Chebyshev’s inequality, that the probability of a creature mistaking the distance to the source approaches zero in the limit as γ grows without bound.⁵ Consequently we introduce a second model, an idealization of Model I, in which the creature *infallibly* determines the distance to the food source on each observation (although not necessarily which direction it lies in). In this model the observation Z_t is just D_t , the actual distance to the source. The *Distance-Certain Model*, Model II, allows full analytical treatment of the asymptotic behavior of the creature under the different strategies of interest.

1.4 Strategies

We focus on three search strategies. The *Information Maximization* (or *Infomax*) strategy seeks to minimize the creature’s uncertainty about the source location, as quantified by the entropy (see Appendix B) of a probability mass function (PMF) representing the possible next source location, given the creatures’ past observations. If multiple locations satisfy this information-greedy algorithm, the creature randomly chooses between the locations with the highest expected food.⁶

For purposes of comparison we introduce a second strategy, the *Likelihood Maximization* (or *Max-Likelihood*) strategy. At each iteration, the creature following this strategy will move to the most likely next location of the source based on its prior observations. As the food distribution peaks at the location of the source, this strategy is equivalent to a greedy or short-term maximization of expected gain. (However, it does not necessarily maximize average expected gain over the long run.)

In addition to the Infomax and Max-Likelihood strategies, we consider a third strategy combining elements of the first two. Under the *modified Maximum Likelihood* (or *Hybrid*) strategy, the creature moves according to the Max-Likelihood strategy when it is *certain* where the source is located, and otherwise it follows the Infomax strategy. That is, the Hybrid creature attempts to minimize its uncertainty about the source until it actually locates it, after which it tries to colocalize with the source even at the risk of subsequently increasing its uncertainty.⁷

We draw a distinction between *survival strategies* (“maximize information”; “maximize the like-

⁵Numerical simulations of this system show that moderately high injection rates lead to very low probabilities of error in estimating the distance to the source (not shown).

⁶Alternatively the creature could chose the *closest* location among those to which the algorithm is otherwise indifferent, rather than choosing several at random.

⁷This strategy could also be called *track and pounce*.

likelihood of colocation”) and *movement algorithms*, the latter being an implementation of a given strategy.⁸ Appendix F provides precise definitions of the algorithm corresponding to each of the three strategies described above. The technical definitions of the Infomax and Max-Likelihood strategies will depend on the notion of a probability mass function encapsulating a creature’s history (Section 2.1); the definition of the Hybrid strategy will depend in addition on the definition of an information state (Section 2.2). These entities are introduced in Section 2.1. While the algorithms are well defined for the full range of values of the source mobility parameter, $0 \leq x \leq 1/2$, they only strictly correspond to implementations of the strategies for which they are named over a narrower range, $1/4 < x < 1/3$ (for further discussion of this point please see Section 3.3.) We will focus our analysis on this narrower range of source mobilities.

Although developed independently, our model system fits comfortably within the closed-loop information/control theoretic framework described in [38]. The state variable subject to control (X , in [38]) is the displacement $S^t - C^t$ between the food source and the creature. The “noise” introduced by the “environment” (\mathcal{E}) corresponds to the random walk performed by the food source, or $S^{t+1} - S^t$. The measurement (\mathcal{A}) is either the molecule count at location C^t or else the observed distance $|S^t - C^t|$. The control apparatus in their system (\mathcal{C}) corresponds to the strategy, along with the internal state of the creature, that together dictate the creature’s next move.

Our model could also be seen as a special case of the general “perception-action loop as a Bayesian network” discussed in ([17], cf. their Figure 1). In contrast with the particular model studied in [17], our model allows an analytic, asymptotic analysis, rather than a computational analysis over a fixed number of time steps; and we consider the consequences of multiple behavioral strategies rather than information maximization alone.

2 Results

2.1 Injector’s Location Given a Creature’s History

The injector’s location at time t , S_t , is a uniform random variable on the ring world (as shown in Lemma F.1). However, the creature can use its history of observations to create a more informative distribution of the injector’s current and (more importantly) next location. We let F_t^- be the probability mass function representing the injector’s location at time t given the creature’s prior history (the history at time $t - 1$). Furthermore, we let F_t^+ be the random PMF of the injector’s location at time t given the creature’s history at time t . (See Appendix D for formal definitions of histories and the PMFs F^- and F^+ .) Note the PMFs F_t^- and F_t^+ are themselves random variables taking values in $\mathcal{P}(\mathbb{Z}_N)$. Given F_t^+ , it is easy to calculate F_{t+1}^- as the convolution of F_t^+ with the injector’s movement algorithm, P :

$$F_{t+1}^- = F_t^+ * P. \quad (2.1)$$

Clearly $F_0^- \equiv 1/N$, the uniform distribution, due to the injector’s starting at a random location with respect to the creature.⁹ Henceforth we will suppress the subscript t ; all random variables

⁸When there is no risk of ambiguity we may refer to strategies and the corresponding algorithms interchangeably.

⁹This statement and others are technically only true “with probability one”. We dispense with this terminology except where necessary for clarity.

will refer to the same time unless noted otherwise.

After a creature makes the observation $Z = z$, its estimate of the source’s location, represented by the PMF F^+ , is updated in a Bayesian fashion (see Appendix E). The update is based on the conditional probability of making the observation Z given the creature’s location (C) and the possible source location $S = l$:

$$F^+(l) = \frac{\mathbb{P}(Z = z|S = l \cap C = c)F^-(l)}{\sum_m \mathbb{P}(Z = z|S = m \cap C = c)F^-(m)}. \quad (2.2)$$

The resulting expression, while somewhat unwieldy in the case of Model I, is nevertheless tractable because of the assumed Poisson nature of the molecule counts [2]. In the Distance-Certain Model, where the observation Z is just the distance D , the right hand side of Equation 2.2 simplifies considerably. The conditional probability of observing a given distance, $\mathbb{P}(D = d|S = l \cap C = c)$, equals one if the distance between l and c is equal to d , and zero otherwise. Moreover, when D is exactly equal to either zero or $N/2$, then Equation 2.2 reduces to

$$F^+(l) = \begin{cases} 1 & l = C + D \\ 0 & \text{otherwise.} \end{cases} \quad (2.3)$$

And when $0 > D > N/2$, Equation 2.2 reduces to

$$F^+(l) = \begin{cases} \frac{F^-(C+D)}{F^-(C+D)+F^-(C-D)} & l = C + D \\ \frac{F^-(C-D)}{F^-(C+D)+F^-(C-D)} & l = C - D \\ 0 & \text{otherwise.} \end{cases} \quad (2.4)$$

Equation 2.3 differs from Equation 2.4 because when D is exactly 0 or $N/2$, then $C + D = C - D$, mod N . Figure 2.1 shows examples of probability mass functions arising in the Distance-Certain Model.

It is biologically implausible to assume that a creature could retain an unlimited list of all of its past locations and observations in order to implement a search strategy. Fortunately, Equations 2.1 and 2.2 demonstrate that with regard to the injector’s location, a lifetime of observations and locations can be captured *exactly* by an appropriate probability mass function. What is more, in both the Discrete Molecule and Distance-Certain models, the PMF required can be updated iteratively in a lossless fashion to incorporate the creature’s latest observation and location. Finally, because the PMF of the injector’s next location (conditioned on a history) can be calculated in a lossless fashion, there is no theoretical limit on the ability of past observations to accurately inform future predictions.

2.2 Coarse Graining: Information States

In Section 1.4 a movement algorithm is defined as a (possibly random) mapping from creature histories to the creature’s next location. We now define a *Markovian movement algorithm* to be any mapping of a creature’s probability mass function F_{t+1}^- (describing the injector’s *next* location, given the creature’s history at time t) to the creature’s next location C_{t+1} . As with a movement algorithm, a *Markovian* movement algorithm may be random or deterministic. Any Markovian

algorithm defines a random process on the state comprising the source, creature, observations, and the PMFs F^- and F^+ . By construction, this random walk satisfies the Markov property, i.e. given its state at time t , its state at time $t + 1$ is conditionally independent of its states at all preceding times $t' < t$. The algorithm corresponding to each of the strategies to be considered (Infomax, Max-Likelihood, and Hybrid) is Markovian, as will become evident in Section 2.4.

In both the Discrete Molecule and Distance-Certain models, the random walk takes place on an infinite state space. For the Distance-Certain model, however, it is possible to define a finite number of coarse-grained *information states* with respect to which the system constitutes a finite dimensional Markov process. In most instances, coarse graining a Markov process on a network produces a random process on a smaller network that no longer satisfies the Markov property. However, for each of the Markov strategies considered, the coarse graining we construct maps the random walk on the original “microscopic” states of the Distance-Certain model to information states in such a way as to preserve the Markovian character of the random walk (as we will show in this Section). Consequently we are able to apply standard results from the theory of Markov processes to obtain analytical results that in turn shed light on the biological merits of the different strategies.¹⁰

The information states not only provide a coarse graining that preserves the Markov property for the system, they also correspond to intuitively appealing equivalence classes that capture biologically salient conditions such as the actual distance between creature and source and the creature’s knowledge about the source. The classification of information states turns on the structure of the probability mass function F^+ reflecting the likelihood that the source is at a given position relative to the creature. We define \mathcal{J}^+ to be the set of PMFs for which the injector’s location is known, or is known to be some definite distance $d > 0$ from the creature’s location. In the Distance-Certain Model the probability that F^+ is in \mathcal{J}^+ is equal to one, by construction. The set \mathcal{J}^+ can be partitioned into the sets \mathcal{J}^{a+} with $a \in \{0, 1, 2, 3\}$ and the set \mathcal{J}^{I+} . For $a \in \{0, 1, 2\}$ the set \mathcal{J}^{a+} is the set of PMFs representing the condition that there is some location, e , such that the injector is equally likely to be at the location(s) $e \pm a$ and has zero probability of being elsewhere. The set \mathcal{J}^{3+} is the set PMFs where the injector is equally likely to be at the locations $e \pm d$ with $d \geq 3$ and has zero probability of being elsewhere. Figure 2.1 illustrates typical PMFs of these states. The set \mathcal{J}^{I+} contains the remaining PMFs in \mathcal{J}^+ , for instance those for which it is certain the injector is a given distance from some location e , but the two possible locations do not have equal probability.

For each probability mass function $f^+ \in \mathcal{J}^+$ representing the likelihood of the source’s current location, there is a corresponding PMF representing the likelihood of its next location. We define \mathcal{J}^- as the set of PMFs that can be generated by Equation 2.1 given $F^+ \in \mathcal{J}^+$; for completeness we also include in \mathcal{J}^- the uniform PMF to account for the initial state of ignorance about the source, F_0^- . We partition the set \mathcal{J}^- into subsets \mathcal{J}^{a-} with $a \in \{0, 1, 2, 3, I\}$ and $\mathcal{J}^{*-} = \{f_u\}$ (the uniform distribution, $f_u \equiv 1/N$). Except for the set \mathcal{J}^{*-} , a PMF f^- is in the set \mathcal{J}^{a-} if and only if there is a PMF f^+ in \mathcal{J}^{a+} such that $f^- = f^+ * P_x$. Figure 2.2 shows typical PMFs of these states. Each such partition corresponds to an information state; we say a creature is *in the information state* $\mathcal{J}^{a\pm}$ when the random PMF F^\pm takes a value f^\pm in the state $\mathcal{J}^{a\pm}$.

¹⁰In the Markov chain literature, an exact coarse-graining is also called a *lumping* of a Markov process.

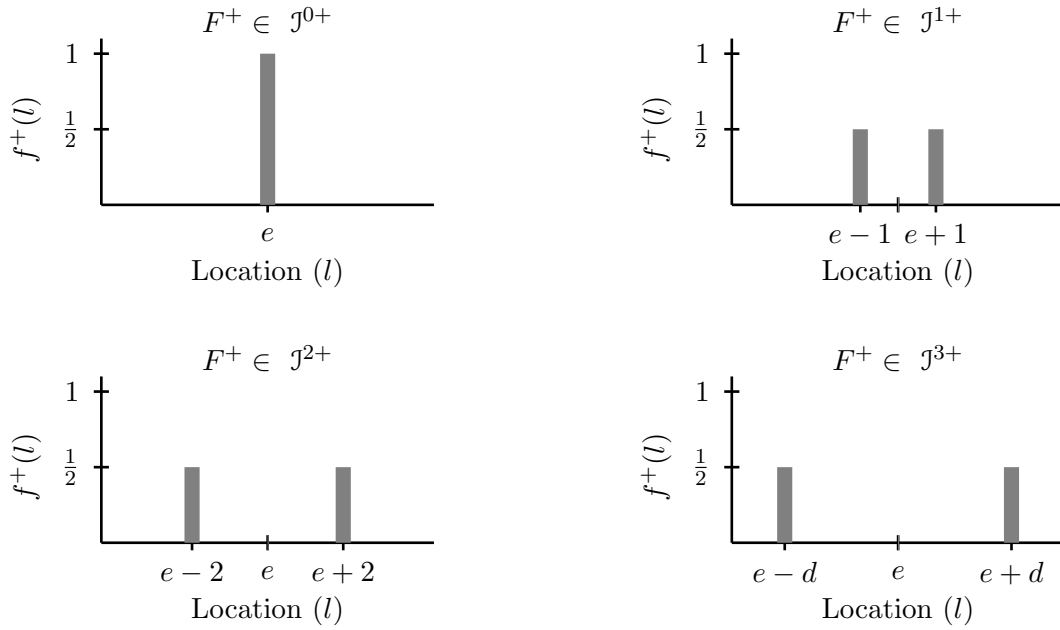


Figure 2.1: The primary information states for the set of probability mass functions \mathcal{J}^+ , representing the conditional probability of the current source location given a creature's history of observations through the present time. The function f^+ is a probability mass function on the ring world \mathbb{Z}_N . For every f^+ in \mathcal{J}^{0+} there is some unique location $0 \leq e \leq N-1$ such that $f^+(e) = 1$ and $f^+(l) = 0$ for all other locations $l \neq e$. For the other information states there are two locations equidistant from some location e where the injector is equally likely to be; furthermore the probability of the injector being at any other location is zero. That is, $f^+(e-d) = f^+(e+d) = 1/2$, and $f^+(l) = 0$ for $l \neq e \pm d$. For f^+ to be in \mathcal{J}^{1+} or \mathcal{J}^{2+} , d must equal one or two, respectively. For f^+ to be in \mathcal{J}^{3+} , d must be greater than two but less than or equal to $N/4$.

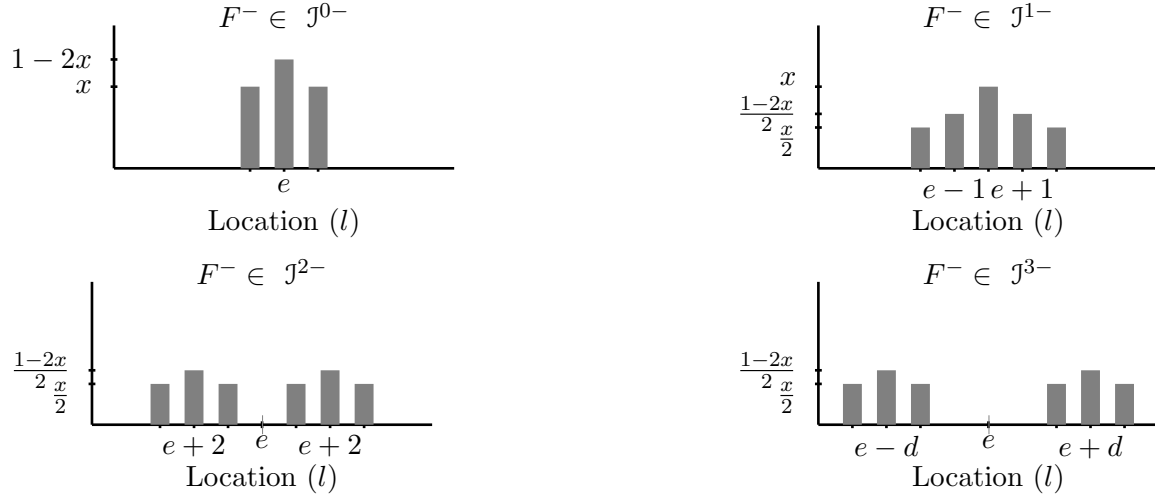


Figure 2.2: Information states following injector movement, and before the next observation. The function f^- is a probability mass function on the ring world \mathbb{Z}_N , representing the likelihood of finding the food source at a given location *after* its subsequent move. Each f^- in \mathcal{J}^- is obtained by convolving a corresponding PMF f^+ in \mathcal{J}^+ with the transition matrix kernel $[\dots, x, (1-2x), x, \dots]$ for the source movement generator P_x . Compare Figure 2.1. Note the shapes of the PMFs depicted are qualitatively correct provided we restrict the range of the source mobility parameter x . For instance a PMF in \mathcal{J}^{1-} takes the values $f^-(e) = x$, $f^-(e \pm 1) = (1-2x)/2$, and $f^-(e \pm 2) = x/2$ for some location e in the ring world, $0 \leq e \leq N-1$. The PMF has the shape shown (increasing monotonically to a central peak at location e) provided $x/2 < (1-2x)/2 < x$, which is equivalent to assuming $1/4 < x < 1/3$. For this figure we set $x = 0.3$.

2.3 Markov Chains on Information States

When the creature is in the information state \mathcal{J}^{0-} (see Figure 2.2) it has a PMF of the form

$$f^-(l) = \begin{cases} x & l = e \pm 1 \\ 1 - 2x & l = e \\ 0 & \text{otherwise.} \end{cases} \quad (2.5)$$

Because in this state the current location of the source is known with certainty, the location of the source at the next time step is known to obey the distribution given in Equation 2.5. Under the Infomax strategy, the creature in information state \mathcal{J}^{0-} must chose to move to location $e + 1$ (or equivalently, $e - 1$) because this move ensures that the creature will know the state with certainty again after the next observation. To see this one need only consider three cases. Given that $S_t = e$ is known, if $C_{t+1} = e - 1$ is chosen then the displacement, $S_{t+1} - C_{t+1}$, can only take one of three values: 0 if the source moves CCW by one step; 1 if the source does not move; 2 if the source moves CW by one step. The Max-Likelihood strategy, in contrast, calls for setting $C_{t+1} = e$, i.e. placing the creature at the next time step at the most likely location of the source. However, with probability $2x$, the source will have moved either CW and CCW, and the creature's observation of the molecule count or the distance, taken at C_{t+1} will leave the new location of the source ambiguous. In this case the goals of minimizing uncertainty and maximizing immediate food gain lead to two different choices. Similar calculations lead to a set of (probabilistic) transitions between the information states for each of the three strategies considered. We can capture the transitions induced by each strategy succinctly in a matrix representation that makes the Markov chain structure of the coarse grained system explicit. Let p_t be the (row) vector representing the probability of being in the initial state (\mathcal{J}^{*-}), or state \mathcal{J}^{a-} , $a = 0, 1, 2, 3$, respectively, on iteration t . Then for each strategy we have a 5×5 matrix M such that $p_{t+1} = p_t M$. Table 2.1 shows the resulting transition matrices for the three strategies considered.

The initial information state is necessarily the maximally uninformed state \mathcal{J}^{*-} with initial PMF $F^- \equiv 1/N$, regardless of the creature's strategy. After the first measurement the creature enters one of the other states \mathcal{J}^{a-} , $0 \leq a \leq 3$. For each transition matrix M there is a subset of information states comprising a minimal absorbing set. It is straightforward to show that for each strategy considered, the Markov chain when restricted to the minimal absorbing set is positive recurrent (see Appendix F.3). Therefore by the Perron-Frobenius theorem there is a unique stationary distribution $\pi \in \mathcal{P}(\mathbb{Z}_N)$ such that $\pi M = \pi$ and $p_t \rightarrow \pi$ as $t \rightarrow \infty$. The stationary distribution may be calculated explicitly for each strategy and for each value of the source mobility, x . The properties of this asymptotic distribution of state occupancies in turn allow us to calculate the average amount of food obtained in the long run for each strategy, as well as the mean uncertainty (entropy) of the source position at each step, given the creature's prior observations.

As an example we will present the calculation of the expected food gained for the simplest case. Suppose, as above, that the creature is in information state \mathcal{J}^{0-} , i.e. at time t the creature knows with certainty that the source is currently at some location $S_t = e$. Recall that the food concentration (Model II) or mean food count (Model I) m steps CW from the source is given by λ_m , $0 \leq m \leq N-1$. If the creature chooses as its next location $C_{t+1} = e + 1$ (or equivalently $e - 1$), then there is a probability x that the creature lands on the source and sees (on average) a quantity λ_0 of food. With probability $1 - 2x$, the source will be at location e (one step away) and so the creature will see an average food count of λ_1 . Finally the creature will see on average λ_2 when the source is at

Infomax	$M_{ITC} = \begin{pmatrix} 0 & \frac{2}{N} & \frac{4}{N} & \frac{4}{N} & \frac{N-10}{N} \\ 0 & \boxed{1} & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix}$
Max-Likelihood	$M_{MLC} = \begin{pmatrix} 0 & \frac{2}{N} & \frac{4}{N} & \frac{4}{N} & \frac{N-10}{N} \\ 0 & \boxed{(1-2\mathbf{x})} & \boxed{2\mathbf{x}} & \boxed{0} & 0 \\ 0 & \mathbf{x} & (1-2\mathbf{x}) & \mathbf{x} & 0 \\ 0 & \boxed{(1-\mathbf{x})} & \mathbf{x} & \boxed{0} & 0 \\ 0 & (1-x) & x & 0 & 0 \end{pmatrix}$
Hybrid	$M_{mMLC} = \begin{pmatrix} 0 & \frac{2}{N} & \frac{4}{N} & \frac{4}{N} & \frac{N-10}{N} \\ 0 & \boxed{(1-2\mathbf{x})} & \boxed{2\mathbf{x}} & 0 & 0 \\ 0 & \boxed{1} & \boxed{0} & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \end{pmatrix}$

Table 2.1: Transition matrices for each strategy. TOP: Information Theory Creature (ITC). CENTER: Maximum Likelihood Creature (MLC). BOTTOM: modified Maximum Likelihood Creature, or Hybrid (mMLC). Each row represents the probability, under the given strategy, of moving from one state (in order: $\mathcal{J}^{*-}, \mathcal{J}^{0-}, \mathcal{J}^{1-}, \mathcal{J}^{2-}, \mathcal{J}^{3-}$) to any of the same five states. Each row sums to unity. For M_{ITC} , there is a minimal absorbing set comprising a single state; for M_{MLC} there is a minimal absorbing set of three states, and for M_{mMLC} (the hybrid strategy) there are two. In each case the absorbing set is indicated by the boxed terms.

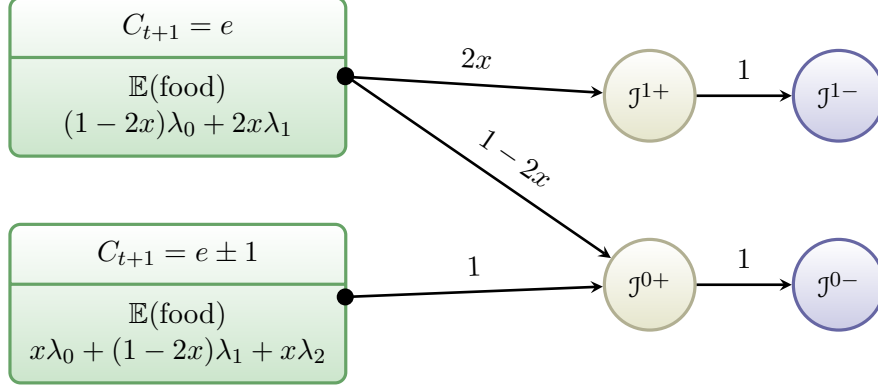


Figure 2.3: Transitions from J^{0-} . This figure states the expected food at each choice for the creature’s next location (C_{t+1}) when in the information state J^{0-} . It also depicts the transition probabilities to other information states based on C_{t+1} . By definition, if the creature is in J^{0+} or J^{1+} at time t the creature will be in J^{0-} or J^{1-} at time $t + 1$ (respectively). Future graphs will not depict the intermediate information states.

$e - 1$ with probability x . Therefore the expected food for location $e + 1$ when in the state J^{0-} is $x\lambda_0 + (1 - 2x)\lambda_1 + x\lambda_2$. The other cases proceed similarly; Figures 2.3, 2.4, 2.5, and 2.6 present the results on transitions between information states and expected food gains for each recurrent information state.

2.4 Strategies Defined by Information States

As defined earlier, a Markov algorithm is any strategy that (perhaps stochastically) maps all possible PMFs to a choice for the creature’s next location. The Infomax, Max-Likelihood, and Hybrid algorithms are depicted in Figure 2.7. It is straightforward to prove (see Lemma F.3) that the Infomax strategy is in fact information optimal in the appropriate sense. For comparison, we may consider the Max-Likelihood strategy, under which the creature always chooses to move to the location at which the injector is most likely to be situated after its next move (with a coin toss in case of a tie). The transitions depicted in Figure 2.7 define the Max-Likelihood strategy provided the parameter x describing the source random walk P_x is between $1/4$ and $1/3$. The Hybrid strategy, also depicted in Figure 2.7, makes the same choices as the Infomax strategy in the states J^{1-} , J^{2-} , and J^{3-} ; however, it makes the same choice as the Max-Likelihood strategy in state J^{0-} .

After two observations, all three strategies enter the absorbing states. In the long term, there is a non-zero probability that the creature is in each absorbing state, and the creature “almost certainly” remains in the absorbing set. These states are highlighted for each strategy in Figure 2.8. As the transition probabilities for these finite, absorbing, coarse-grained states are exact, it is possible to calculate the expected time the creature will spend in each state (see Appendix G).

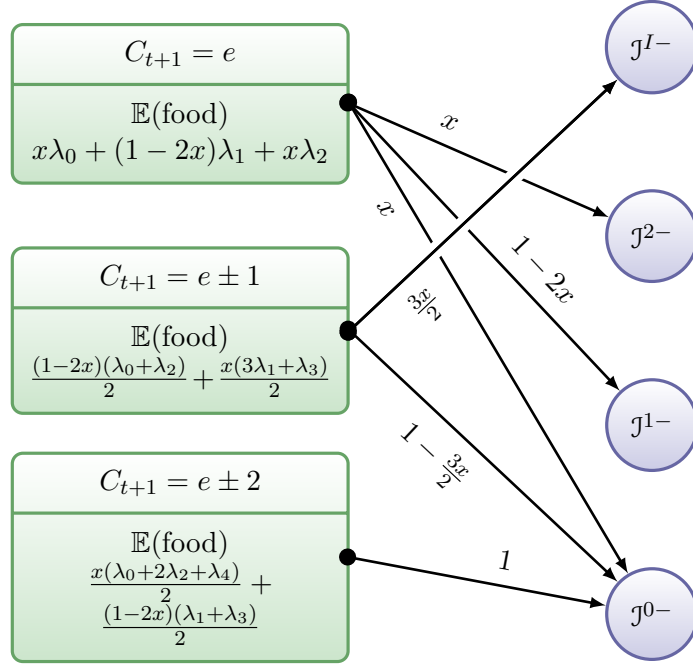


Figure 2.4: Transitions from J^{1-} . This figure states the expected food at each choice for the creature's next location (C_{t+1}) when in the information state J^{1-} . It also depicts the transition probabilities to other information states based on C_{t+1} .

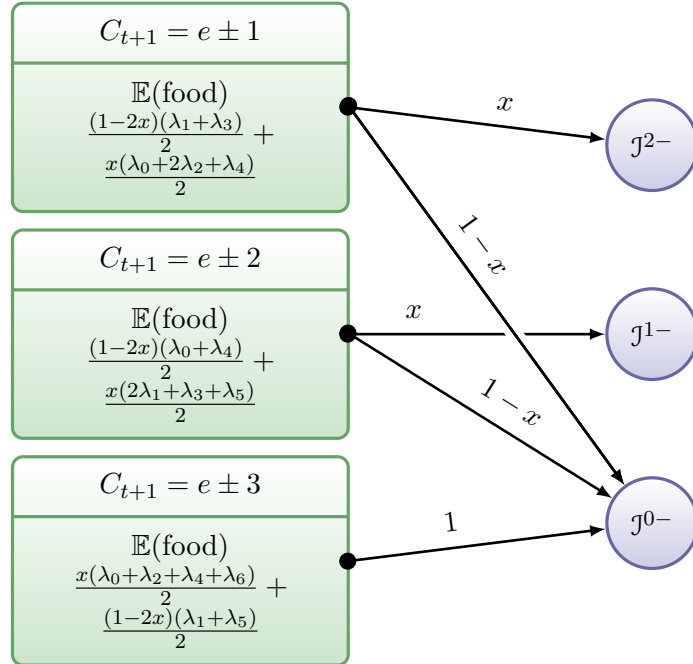


Figure 2.5: Transitions from J^{2-} . This figure states the expected food at each choice for the creature's next location (C_{t+1}) when in the information state J^{2-} . It also depicts the transition probabilities to other information states based on C_{t+1} .

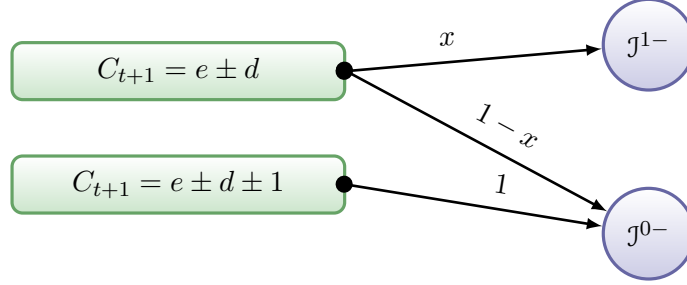


Figure 2.6: Transitions from J^{3-} . The expected food in this state is a function of d ; the creature strategies do not spend more than one turn in this state so the effect on the expected food over time is negligible. It does depict the transition probabilities to other information states based on C_{t+1} .

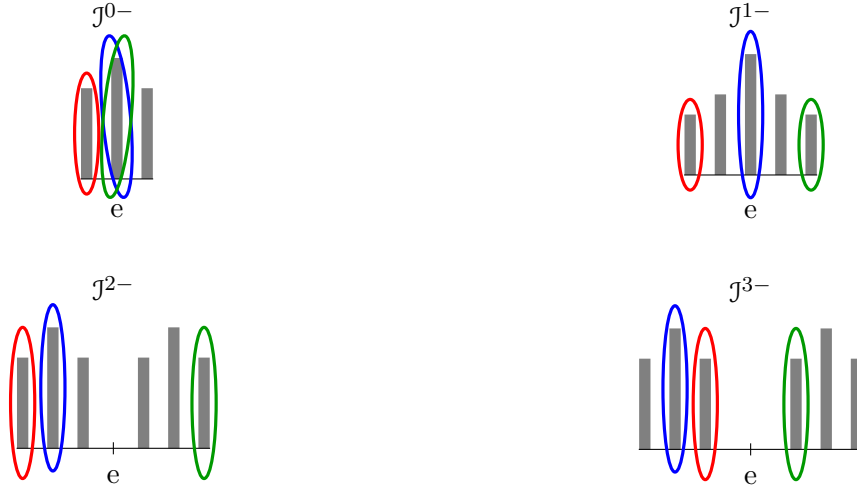
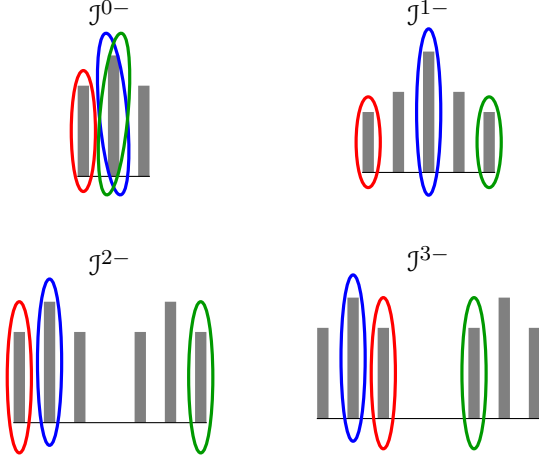
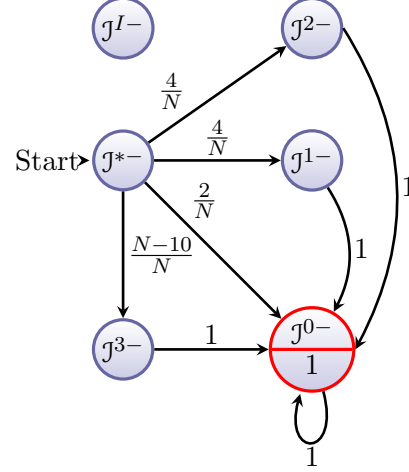


Figure 2.7: The **Infomax** (red), **Max-Likelihood** (blue), and **Hybrid** (green) strategies are defined by the choice of location given F^- is in a particular information state. In each panel above, a colored oval indicates the location to which a creature following the given strategy would move at time $t + 1$, starting from the corresponding information state at time t . For example, if the creature is certain that the source is currently at a given location (say e), *i.e.* the creature is in information state J^{0+} (Top Left panel), then under the **Max-Likelihood** and **Hybrid** strategies the creature would move to location e , the most likely location for the source to be at the next step (provided the probability of the source not moving is greater than $1/3$ or equivalently $x < 1/3$). But the **Infomax** strategy would move to $e \pm 1$ with equal likelihood, because this location minimizes the expected uncertainty about the source's location at the next time step. The remaining panels are interpreted similarly. In the information state J^{1-} (Top Right panel), the injector is more likely to be at the location e than at the locations $e \pm 1$ if and only if $x > 1/4$. Therefore the Max-Likelihood *algorithm* as depicted in this figure is consistent with the likelihood maximization *strategy* only for $1/4 < x < 1/3$. By symmetry, if a strategy would choose the location $e + m$ with $m \in \{0, \dots, N - 1\}$, the transition probabilities to different states and the expected food would be the same at location $e - m$. Therefore, the set of strategies that choose $e + m$ with probability $p \in [0, 1]$ and $e - m$ with probability $1 - p$ in a given information state make equivalent choices for C_{t+1} with respect to the information-state coarse graining.

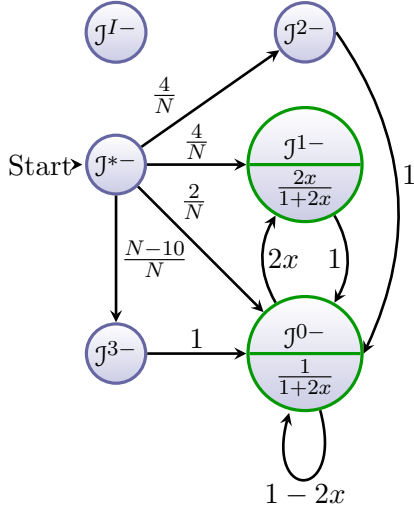
Strategy Definitions



Information Theory Creature



Hybrid Creature



Maximum Likelihood Creature

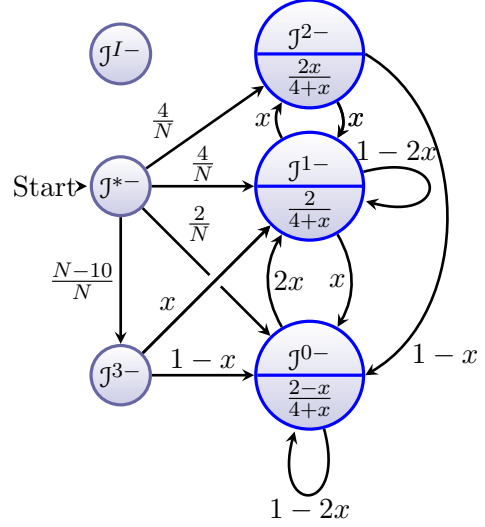


Figure 2.8: Transition Graphs for each of the strategies. Top left panel recapitulates Figure 2.7. Other panels depict transitions between information states for each strategy as labeled. Each creature begins in state j^{*-} . Arrow labels indicate transition probabilities. Large colored circles denote positive recurrent states for each strategy, with expressions provided for the equilibrium probability of occupying the state. Note the sum of transition probabilities on each set of arrows pointing away from a given node sum to one. For example, the creature always starts in state j^{*-} , the uniform distribution, representing complete ignorance of the injector's location. From this state it moves to state j^{2-} with probability $4/N$, to state j^{1-} with probability $4/N$, to state j^{0-} with probability $2/N$, and to state j^{3-} with probability $(N-10)/N$. These transitions depend only on making the first observation, and are the same for each strategy. Subsequent transition probabilities differ according to the strategy employed, which determines how the creature moves relative to the source.

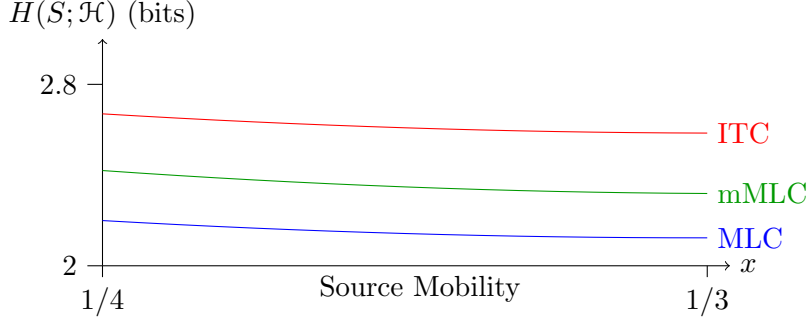


Figure 2.9: The expected information for each strategy as a function of x , the source mobility parameter. Larger values of x correspond to greater unpredictability of future source locations. Each curve represents the expected decrease in uncertainty about the source position when conditioned on the creature’s history and observations, compared with the underlying uncertainty in the absence of any observations ($\log_2(N)$ bits), for the different strategies considered. As expected, the movement strategy based on information maximization (ITC, Information Theory Creature) dominates the maximum likelihood based strategy (MLC, Maximum Likelihood Creature) and the hybrid strategy (mMLC, modified Maximum Likelihood Creature) over the entire range of x considered. Also as expected, the information about the source position relative to the creature’s position that is gained through observation decreases as the source’s inherent unpredictability increases, *i.e.* all three curves have strictly negative slope.

2.5 Strategies Compared

The information a creature has about the injector’s location given an information state is calculated in Lemma F.3. Using the expected time spent in each state for a given strategy¹¹ the expected information for each strategy is shown in Figure 2.9. This figure shows the strict ranking of information theoretic performance between the strategies; namely, maximum likelihood has the least, the hybrid strategy is in the middle, and Infomax is the best. This result, which is not unexpected, is easily explained by examining the fraction of time a creature spends in a given information state. The uncertainty about the source’s location as quantified by the entropy of the source location distribution (see Table G) is strictly increasing as a function of both x and the information state index k . The creature in information state \mathcal{J}^{0+} knows exactly the current location of the food source; creatures in \mathcal{J}^{1+} and \mathcal{J}^{2+} have progressively more uncertainty. The information-seeking creature is able to remain in state \mathcal{J}^{0+} on every time step. In Lemma F.3 we prove that the Infomax strategy is information optimal, in the sense that the creature following the Infomax strategy always has the least uncertainty, *i.e.* the most information, about the source’s location. For all three strategies as the source mobility x increases, the creature’s uncertainty about the source location gradually grows, as seen in the decrease of the mutual information in Figure 2.9. Table 2.2 shows the fraction of time spent by each strategy in each information state; Figure 2.10 displays this comparison graphically.

Figure 2.11 shows the expected food gain for each strategy. Unlike the expected information, *the expected food shows no consistent ranking between strategies*. For some choice of source mobility

¹¹The fraction of time a creature spends in a given information state is exactly the equilibrium probability distribution, derived in Appendix F.3. See also Figure 2.8, Table 2.2 and Figure 2.10).

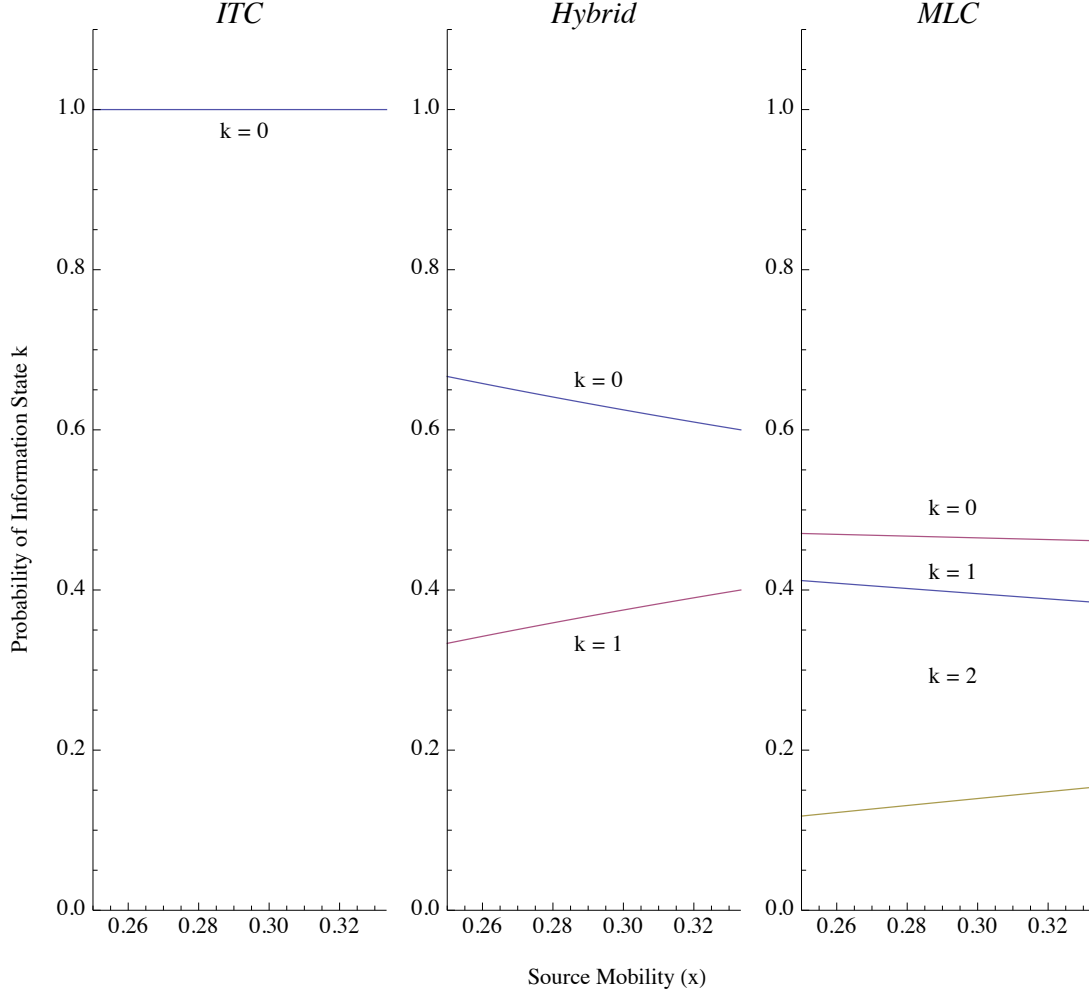


Figure 2.10: Fraction of time spent in the k^{th} information state for each strategy, as a function of the source mobility x . In state \mathcal{I}^{k+} the food source is known to be at position $e \pm k$ for some location e . *ITC*, Information Theory Creature. The creature following this algorithm remains in information state zero for all time after an initial transient. *Hybrid*, modified Maximum Likelihood Creature. The creature following this algorithm asymptotically spends all of its time in information states zero and one. The fraction of time spent in the lowest entropy state (state zero) decreases as the source mobility increases. *MLC*, Maximum Likelihood Creature. The creature following this algorithm asymptotically spends its time in information states zero, one and two. Compare Table 2.2.

k	ITC	$x \rightarrow 1/4$	Hybrid	$x \rightarrow 1/3$	$x \rightarrow 1/4$	MLC	$x \rightarrow 1/3$
0	1	$\frac{2}{3} \leq$	$\frac{1}{1+2x}$	$\leq \frac{3}{5}$	$\frac{7}{17} \approx 0.41 \leq$	$\frac{2-x}{4+x}$	$\leq \frac{5}{13} \approx 0.38$
1	0	$\frac{1}{3} \leq$	$\frac{2x}{1+2x}$	$\leq \frac{2}{5}$	$\frac{8}{17} \approx 0.47 \leq$	$\frac{2}{4+x}$	$\leq \frac{6}{13} \approx 0.46$
2	0		0		$\frac{2}{17} \approx 0.12 \leq$	$\frac{2x}{4+x}$	$\leq \frac{2}{13} \approx 0.15$
3+	0		0			0	

Table 2.2: Fraction of time spent in each information state for each strategy. Each row indicates the fraction of time π_k the creature spends in the corresponding information state as a function of x , and the lower and upper limits of π_k viewed as a function of x . The Information Theory Creature (ITC) remains permanently in the maximally informed state, $k = 0$. The Hybrid creature (mMLC) spends twice as much time in state 0 when x is small as it does in the less informed state 1. As x increases from $x \approx 1/4$ to $x \approx 1/3$, the creature spends only 50% as much time in state 0 as in state 1. The pure MLC creature always spends *more* time in state 0 than in state 1, and an increasing amount of time in state 2 as x increases. Compare Figure 2.10.

in the range $1/4 < x < 1/3$ each strategy outperforms the other two. In particular, the infomax strategy provides a better average food intake than either the maximum likelihood or the hybrid strategy, when the source is the most unpredictable ($x \lesssim 1/3$), but fails to provide the best performance as the next source location becomes easier to predict, given the current location.

It is tempting to explain this result by noting that when $x \lesssim 1/3$, the food source itself generates an ensemble of trajectories that has a higher entropy generation rate than when x is smaller. Perhaps these more “information rich” trajectories somehow cause a creature focused on *information about* the food source to outperform those focused only on *colocation with* the source. Unfortunately such qualitative reasoning would fail to explain why, for intermediate values of the source mobility, it is the maximum likelihood creature that gains more food, not the hybrid that mixes the ML and IT strategies. Also puzzling is the trend of the ITC curve. It is not surprising that the performance of the MLC and mMLC strategies decreases as the unpredictability of the source increases, but why should the ITC performance curve have a *positive* slope?

To understand more fully how “information” and “utility” intertwine to determine the performance of each strategy, we may consider the transition probabilities for the Markovian random walk on the set of states *jointly* describing the creature’s information state and its actual distance from the food source. In a sense this takes us a step backwards as we “un-coarse-grain” the system. With the benefit of hindsight, we take advantage of the fact that for each strategy, the underlying system admits a Markovian coarse graining in terms of the information states alone. Equations (H.4-H.6) in Appendix H provides the transition probabilities for joint (distance, information) states for each strategy. Examination of these transition probabilities shows that two of the strategies (ITC and MLC) admit an exact coarse graining in terms of the distance to the source. Table 2.3 provides the resulting distribution of time spent at each distance from the source, for each strategy; these results are plotted in Figure 2.12.

The expected food obtained under a given strategy is assumed to be the weighted average of the food available at each distance, weighted by the probability of being that distance from the source. The food distribution for model II, as derived in Appendix B (see also Figure 1.3) is strictly

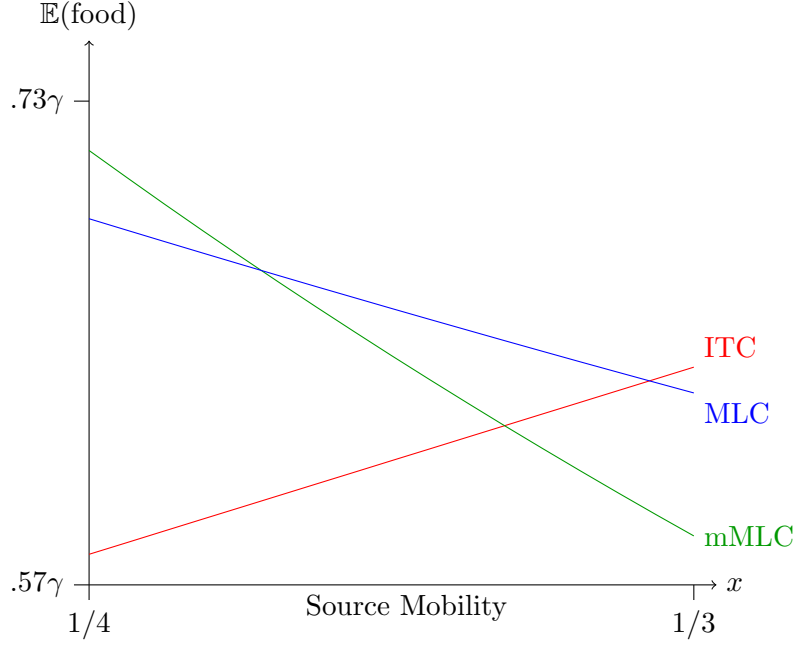


Figure 2.11: The expected food for each strategy as a function of the source mobility parameter x . (ITC, Information Theory Creature. MLC, Maximum Likelihood Creature. mMLC, modified Maximum Likelihood Creature, or Hybrid.) Expected food is calculated as described in Section 2.3. Each strategy is superior to the other two (in terms of long-term food gain) for some range of source mobility. Over the range considered ($1/4 < x < 1/3$), the infomax strategy dominates only in a narrow range corresponding to relatively random motion of the source. The maximum likelihood based strategy dominates over a broad range of mobilities, but for the lowest mobilities the hybrid strategy is the most successful. Compare Figure 2.9.

Distance	ITC	Hybrid	MLC
0	x	$(1-x)^2/(1+2x)$	$(2-2x)/(4+x)$
1	$1-2x$	$(3x-2x^2)/(1+2x)$	$2/(4+x)$
2	x	$2x^2/(1+2x)$	$2x/(4+x)$
3	0	$(x-2x^2)/(1+2x)$	$x^2/(4+x)$
4	0	$x^2/(1+2x)$	$(x-2x^2)/(4+x)$
5	0	0	$x^2/(4+x)$
≥ 6	0	0	0

Table 2.3: Fraction of time spent by each creature a given absolute distance from the food source, as a function of the mobility parameter x . ITC, Information Theory Creature; Hybrid, modified Maximum Likelihood Creature; MLC; Maximum Likelihood Creature. Compare Figure 2.12.

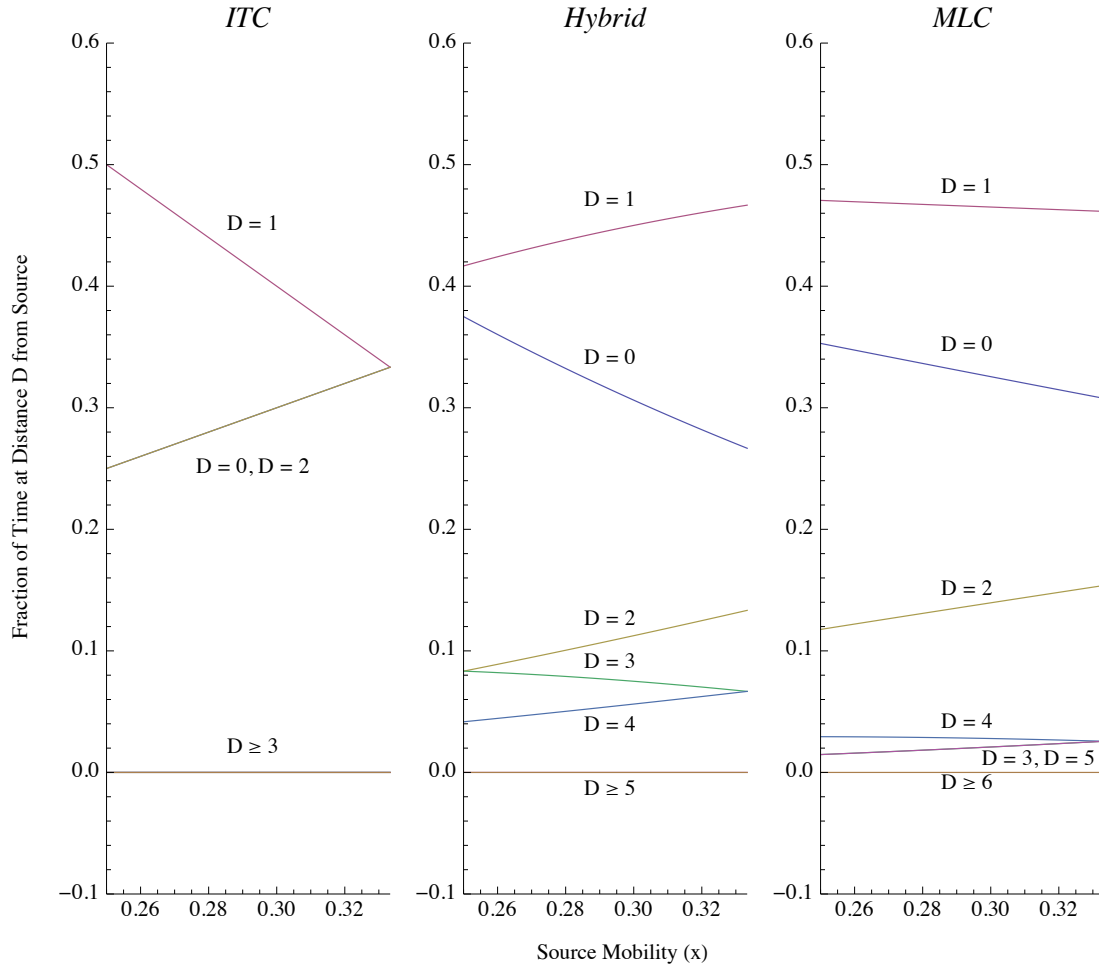


Figure 2.12: Fraction of time spent by the creature at a given absolute distance from the food source for each strategy, as a function of the mobility parameter x . *ITC*, Information Theory Creature. Under this algorithm, the creature spends equal amounts of time at distance $D = 0$ and $D = 2$ from the source. The amount of time at the highest food concentration ($D = 0$) increases as the mobility parameter x increases, while the amount of time spent one step away from the source ($D = 1$) decreases. The probability of the ITC being a distance $D \geq 3$ from the source is zero, after the initial transient. *Hybrid*, modified Maximum Likelihood Creature. As x increases, the mMLC spends more time at distance $D = 1$ and less time at distance $D = 0$. The probability of the mMLC being a distance $D \geq 5$ from the source is zero, after the initial transient. *MLC*, Maximum Likelihood Creature. As x increases the MLC spends slightly less time at distances $D = 1$ and $D = 0$, and more time at distance $D = 2$. The probability of the MLC being a distance $D \geq 6$ from the source is zero, after the initial transient. Compare Table 2.3.

Type	$F(d), 0 \leq d \leq 5$
Concave Up	$(5 - d)^2$
Concave Down	$25 - d^2$
Linear	$25 - 5d$
Delta Function	$25 \delta(d)$

Table 2.4: Examples of different shapes of food distribution functions: concave up, concave down, linear and delta function. Each function has a nominal peak value of $F(0) = 25$; each is taken to be zero for $d \geq 5$. Dependence of expected food gained for each strategy for each function shown is plotted in Figure 2.13.

decreasing in d . Therefore the best conceivable performance would be for the creature to remain at zero distance from the source at all times. The unpredictability of the source motion makes this impossible, and consequently the strategies vary in the fraction of time spent at distance $D = 0$, as a function of the source mobility parameter x . As x increases, both the Max-Likelihood and the Hybrid creatures spend less and less time colocalized with the source, while probability of $D = 0$ increases for the Infomax strategy.

The increase of utility for the ITC with increasing x therefore has a simple, intuitive explanation. The ITC acts in a way that takes into account the actual amount of food to be had at each destination only in a secondary fashion. Aiming directly for the source’s last location will lead, with probability $2x$, to a one bit increase in uncertainty about the source’s location on the subsequent time step. Therefore the IT creature will always choose to miss the most likely location by a single step. It takes the second best position, from the point of view of food collection, in order to have the best position for continued certainty about the source’s trajectory. This strategy is a poor choice on its face, because the only way the creature will collect better than the second highest food amount is if it gets “lucky” and the next move of the source happens to take it to the creature’s location by chance. The probability of colocation if the creature moves to a location one step to the right or left of the source’s last known position is exactly the source mobility x . Hence as x increases the probability of collocation for the IT creature grows rather than decreases.

Quantitatively, the expected food obtained by following the infomax strategy is straightforward to obtain. If the food obtained a distance d from the source is λ_d , then the expected food for the ITC may be written in terms of the second difference (discrete second derivative) of the food distribution:

$$\mathbb{E}_{\text{ITC}}(\text{food}) = x\lambda_0 + (1 - 2x)\lambda_1 + x\lambda_2 \quad (2.6)$$

$$= \lambda_1 + x(\lambda_0 - 2\lambda_1 + \lambda_2) \quad (2.7)$$

Therefore the increase in the expected food gained under the infomax strategy directly reflects the upwards concavity of the food distribution curve as shown in Figure 1.3.

To determine the effect of changing the shape of the food distribution curve about the source, we considered several generic food distribution functions $F(d)$, given in Table 2.4. Three distributions are chosen representing “concave up”, “concave down”, or linear. Each has constant second difference over the relevant distance range, respectively positive, negative, or zero. Each is monotonically decreasing as a function of distance, with a peak at $d = 0$. Because the recurrent states only occupy

distances in the range $0 \leq d \leq 5$, we set our trial food distribution functions to zero for $d \geq 5$. We also consider an all or none dependence of food gained on distance to the source (“delta function” in Table 2.4). In the latter case food *signaling* is decoupled from food *obtained*. Here it is assumed that the creature can determine the distance to the source by detecting the local concentration of a signal (visually or *via* olfaction, for instance), but only captures the food through colocation. This situation would be familiar to any predator.

Figure 2.13 illustrates the results. As anticipated, the slope of the expected food gain under the infomax strategy has the same sign as the second difference of the food distribution. More surprisingly, the topology of the relative rankings of strategies as a function of source mobility appears to be relatively insensitive to the concavity of the food distribution. For the four cases examined (Table 2.4) the rankings appeared in the same order, with the hybrid strategy dominating the others for the lowest source mobilities, the maximum likelihood strategy dominating for intermediate values and the ITC dominating for the highest source mobilities.¹² The combined plots of strategy performance show the same “triangle” present for the distribution considered originally, compare Figure 2.13 and Figure 2.11.¹³

2.6 Distance Traveled

The distance traveled in pursuit of a strategy is another factor that might impact an organism’s survival. While detailed consideration of metabolic constraints would go beyond the scope of this paper, the models considered do lead to simple analytic expressions for two quantities of potential relevance in this context. In Appendix I we indicate how one calculates the expected distance moved per time step by a creature following each of the strategies, as well as the probability of movement per time step (as opposed to the probability of remaining motionless). The average distance moved could be related directly or indirectly to a metabolic movement cost. The probability of movement could also be related to a cost for initiating movement, such as the effort a foraging bird might have to exert in order to ascend to a given altitude, after which soaring arbitrary distances might incur relatively little additional metabolic cost, or to the differential risk of predation during a motile *versus* a sessile phase of activity.

As Table 2.5 indicates, the average distance moved per time step by a creature following the Max-Likelihood strategy is strictly less than that moved by the Infomax creature, which in turn is less than the mean distance moved by the Hybrid creature. Similarly, the probability of movement (the probability that $C_{t+1} - C_t \neq 0$) is least under the MLC strategy and greatest under the Hybrid strategy. Regardless of the detailed nature of the movement cost, the maximum likelihood algorithm would have the greatest advantage in terms of movement efficiency. If resource storage or resource utilization were limiting factors, this might lead to selection against the pure information maximization strategy in favor of the maximum likelihood strategy. Likewise, exposure to predation could also impose a selective cost that might further favor the ML creature over the ITC and especially the Hybrid creature. However it should be noted that as shown in Figure 2.14 the

¹²The topology of relative rankings is not *universally* invariant with respect to the food distribution, however. For instance, choosing $\lambda_0 = 5, \lambda_1 = 4, \lambda_2 = 1, \lambda_{k \geq 3} = 0$ results in $\mathbb{E}_{ITC}(\text{food}) < \mathbb{E}_{MLC}(\text{food})$ when $x = 1/3$.

¹³In order to see the full topology of the curves, we plot the expected food over a range of mobility values extending to lower mobility than the $x = 1/4$ cutoff used in the rest of the paper. For a discussion of the applicability of the different movement algorithms outside the range $1/4 < x < 1/3$, see Section 3.3.

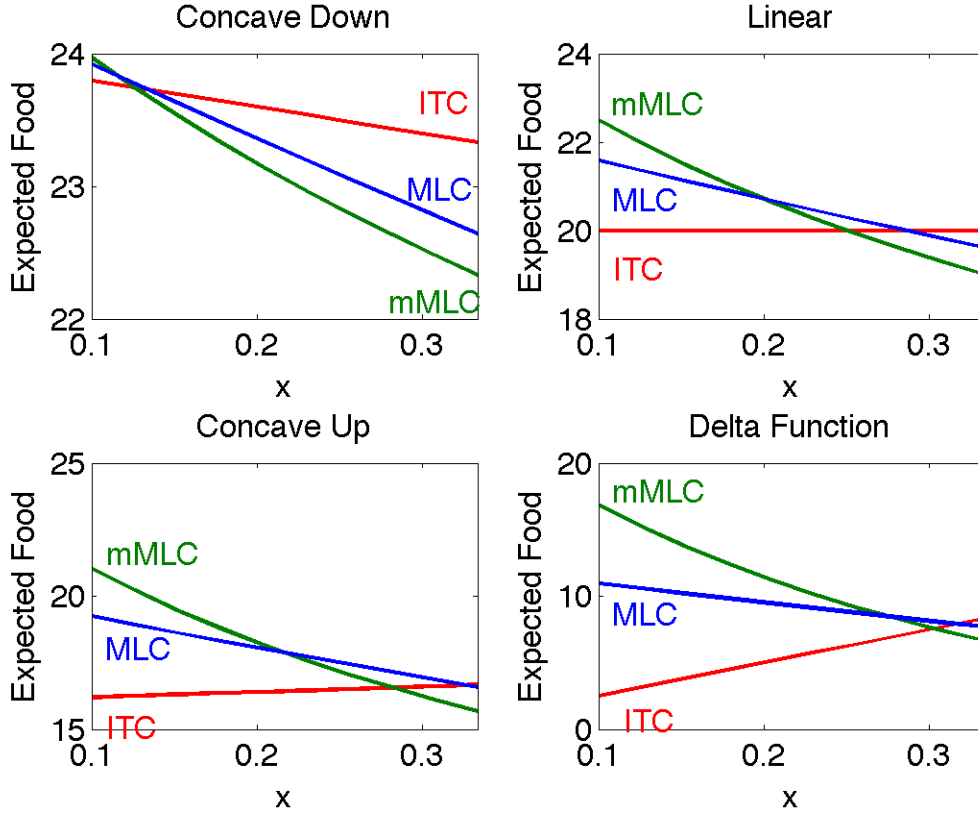


Figure 2.13: Persistence of the ranking vis-a-vis expected food gathered across different food distribution shapes. Each plot shows expected food gained, in nominal units, versus source mobility on an expanded scale $1/10 < x < 1/3$. Food gained as a function of distance from the source, $F(d)$, was taken to have constant second difference in three generic cases. The nominal food function $F(d)$ was set to zero for $d \geq 5$ for this comparison. Concave Down: $F(d) = 25 - d^2$. Linear: $F(d) = 25 - 5d$. Concave Up: $F(d) = (5 - d)^2$. Delta Function: $F(d) = 25$ if $d = 0$, otherwise $F = 0$. In each case, there is an upper range in which the Infomax movement algorithm (red curve) dominates, a middle range in which the Max-Likelihood movement algorithm (blue curve) dominates, and a lower range in which the Hybrid movement algorithm (green curve) dominates. In each case the curve for the Infomax algorithm is a straight line with slope either positive (for “Concave Up” and “Delta Function” food distributions), negative (for “Concave Down” food distribution) or zero (for “Linear” food distribution).

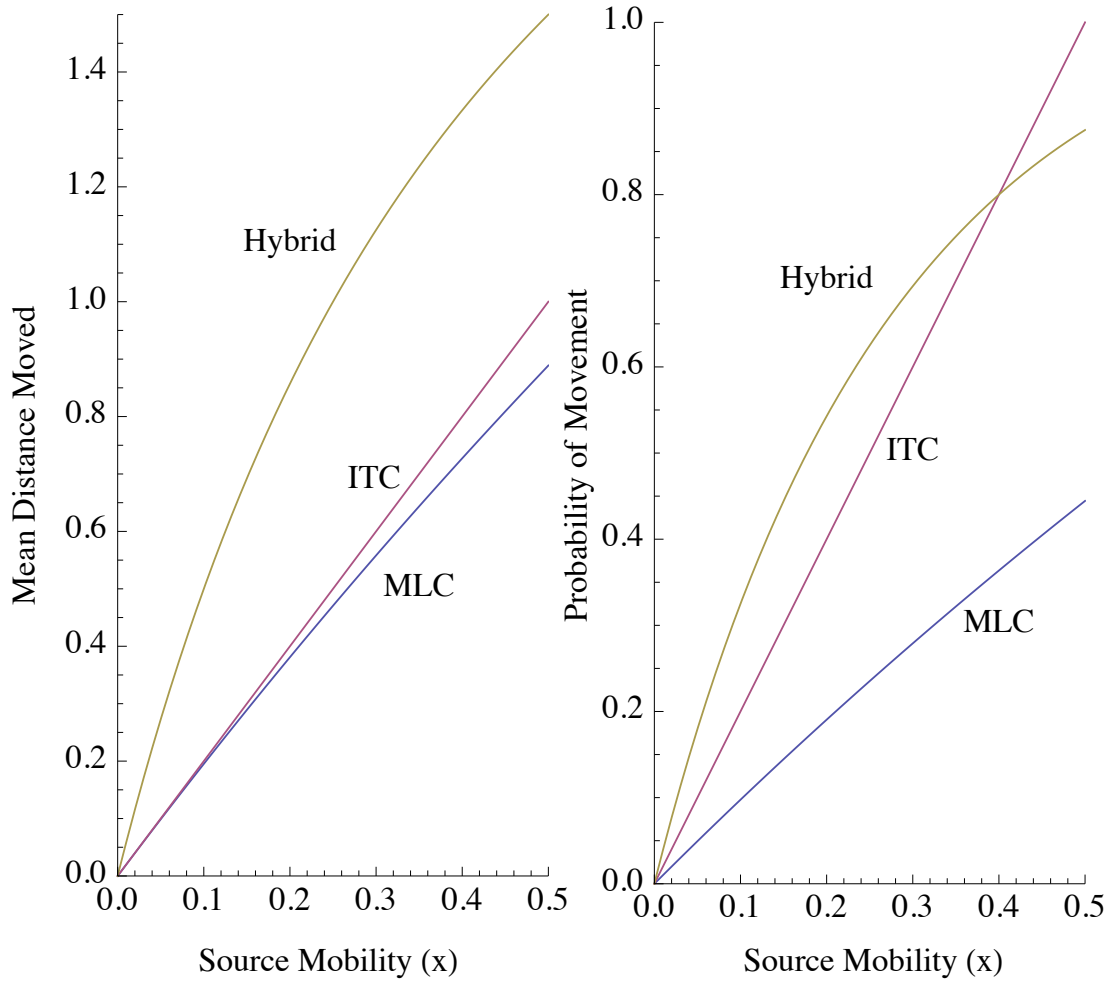


Figure 2.14: Mean distance traveled and probability of moving per time step. MLC: Maximum Likelihood Creature, ITC: Information Theory Creature, mMLC: modified Maximum Likelihood Creature. **Left Panel:** the average distance moved per time step as a function of the mobility parameter x . **Right Panel:** the probability of moving (as opposed to remaining in place) on each time step as a function of the mobility parameter x . Over the range $1/4 < x < 1/3$ the movement algorithms strictly correspond to the movement strategies. Within this range, the MLC moves less frequently and moves a shorter distance on average than the ITC or the Hybrid creature. However, the mean distance moved by the MLC is only slightly less than that moved by the ITC. Compare Table 2.5.

	MLC		ITC		mMLC	
$\mathbb{E}[C_{t+1} - C_t]$	$\frac{8x}{4+x}$	$<$	$2x$	$<$	$\frac{6x}{1+2x}$	(for $0 < x \leq 1/2$)
$\Pr[C_{t+1} \neq C_t]$	$\frac{4x}{4+x}$	$<$	$2x$	$<$	$\frac{x(4-x)}{1+2x}$	(for $0 < x < 0.4$)

Table 2.5: Mean distance traveled and probability of moving per time step. MLC: Maximum Likelihood Creature, ITC: Information Theory Creature, mMLC: modified Maximum Likelihood Creature. The mean, or expected distance traveled in one time step is $\mathbb{E}[|C_{t+1} - C_t|]$. The probability of moving on a given time step is the probability that C_t and C_{t+1} take different values, *i.e.* $\Pr[C_{t+1} \neq C_t]$. The last column shows the range of validity of the inequalities. For the second row, the inequality $\text{MLC} < \text{ITC}$ holds for all $x > 0$, but the inequality $\text{ITC} < \text{mMLC}$ reverses when $x > 2/5$. Compare Figure 2.14.

advantage enjoyed by the MLC in terms of mean distance traveled is very slight compared to the ITC, whereas the absolute difference in the probability of moving is more significant. Together our results reinforce the conclusion that while information efficiency is surely one important factor in the evolution of adaptive behavior, the significance of information efficiency relative to other considerations is nuanced and context dependent.

3 Discussion and Conclusions

3.1 Summary & Conclusions

We developed an analytically tractable model for a particular instance of a perception-action loop: a creature searching for a wandering food source confined to a one-dimensional ring world. The model encompassed the statistical structure of the possible observations, the effects of the creature’s actions on that structure, and the creature’s strategic decision making. Although the underlying model took the form of a Markov process on an infinite dimensional state space, we successfully devised an exact coarse graining that reduced the model to a Markov process on a finite number of “information states”. This technique allowed us to make quantitative comparisons between the performance of an information-theoretically optimal strategy with other candidate search strategies.

From the results shown in Figures 2.9 and 2.11 in Section 2.5 we conclude for the simple Distance-Certain model system (Model II) that:

1. Information optimal search does not necessarily optimize utility.
2. The rank ordering of strategies by information performance does not predict their ordering by expected food obtained, *i.e.* there is no simple tradeoff between information and utility.
3. The relative advantage of different strategies depends on the statistical structure of the environment, in particular the variability of motion of the source:

- (a) The likelihood maximization strategy outperformed the other two strategies considered (in terms of utility) for a wide range of source mobility.
- (b) For the lowest source mobilities considered ($x \gtrsim 1/4$), the hybrid strategy was superior to the other two.
- (c) The information maximization strategy dominated the other two in a narrow range at the highest source mobility considered ($x \lesssim 1/3$).

Consequently we conclude that the strong hypothesis, “behavioral optimality implies information optimality” is false, at least in this model system.

More broadly, our results suggest that while information efficiency is surely one important factor in the evolution of adaptive behavior, the significance of information efficiency relative to other factors such as expected food gain or expected distance traveled is complex, and needs to be defined through a careful analysis of the organism and its environment. Rather than proposing to replace the strong hypothesis of information efficiency with an equally sweeping general hypothesis, the model developed here points the way towards a more systematic analysis of the interplay between environmental structure and uncertainties, a creature’s needs and means of satisfying them, and the various senses in which it can do so efficiently.

The difficulties inherent in analyzing systems in which the dependencies of action on sensation and sensation on action form an apparently intractable closed loop are well established. In addition to establishing a counterexample against the strong hypothesis, our work also leads to a positive conclusion: that despite these difficulties, one may make progress in understanding the interplay of information processing and utility for survival in behavioral systems by embedding a creature and its world together in a combined Markov process. As we point out in Section 3.5 below, this approach can lead to numerous extensions of the present work.

3.2 Limitations of the Model

How do the simplifications required to obtain a tractable model limit the scope of our conclusions?

The simple model analyzed here incorporates observation, action, and the effects of action on future observations in a nontrivial way to form a closed action-perception cycle. The key aspect of the model we construct is that the system comprising the “creature” and the “world” together constitute a discrete state, discrete time Markov chain amenable to analysis using standard probabilistic tools. One may readily conceive of more elaborate models that include additional elements such as: the cost of memory and computation required to implement a given strategy, imperfect observations, imperfect knowledge of the world, interaction of more than one signal or more than one type of nutrient, constraints such as mortality after excessively long periods without food, goals such as accumulating a sufficient intake of food to achieve reproductive success, and metabolic or predation related movement costs.

Would any of these additional considerations weaken the conclusion that the model creatures considered here establish a counterexample to the strong hypothesis? On the contrary. In the model

used here, the information maximization creature has been given all the possible advantages: perfect information about the operating conditions of the world, perfect measurement capability, exact representation of its internal state (probability distribution for the source location), the ability to move arbitrary distances without cost. If it does not win the competition with all the winds in its favor, it should not stand a chance under more realistic conditions either. In particular, the addition of memory or computational costs for maintaining the internal state should not affect one creature more than another, as all three strategies are based on the same internal representation of the source probability distribution. Removing the assumption that the IT creature has perfect knowledge of the world (aside from the source location) and perfect measurement capability would only weaken its performance. The ITC does not on average move shorter distances or move less frequently than the MLC. It is likely that more realistic models should reinforce our conclusions.

We note the significance of the Bayesian update and convolution for predicting the injector's location. Although we assumed a particular geometry for the world (a ring) and a symmetric random walk for the source, our framework would work equally well for a discrete world of any shape where observations (of any kind) are conditionally independent of prior observations given a source's location S and creature's location C . For instance, we could consider 2D or 3D worlds, and observations governed by distributions other than Poisson. We assumed that the transition matrix P governing movement of the source had translational symmetry, *i.e.* that the *increments* of the injector's location be independent of the injector's location. However, this assumption could be relaxed without changing the general conclusions of the study. For example, the injector could have landmarks that influence its movement, provided the creature knew their location and influence. In this situation we would still have a Markov process with a transition matrix P . The mapping from PMF for the current location F_t^+ to the PMF of the new location F_{t+1}^- would be given by matrix multiplication: $F_{t+1}^- = P \cdot F_t^+$. In this case, the world would not need to be symmetric.

3.3 Survival Strategies *versus* Movement Algorithms

The movement algorithms are defined in terms of destination given an information state \mathcal{J}^- . The algorithms specified in Figure 2.7 remain well defined over the range $0 \leq x \leq 1/2$. However the qualitative properties satisfied by the movement algorithms are only consistent with the given survival strategies (information maximization, likelihood maximization, track and pounce) over narrower ranges of x . The information maximization *strategy* (first maximize information; then maximize expected food among equally informative locations) is consistent with the ITC movement algorithm for all $0 < x \leq 1/2$. For $x \equiv 0$, however, the infomax strategy chooses to colocate with the source following the first observation. The movement algorithm for the MLC as defined in Figure 2.7 is consistent with likelihood maximization only when $1/4 < x < 1/3$. For $0 < x \leq 1/4$, the central location is no longer the most likely location in information state \mathcal{J}^{1-} . Consequently a strategy of strict likelihood maximization will no longer remain within information states $\mathcal{J}^{0-} - \mathcal{J}^{2-}$ and the resulting Markov chain will no longer be absorbed by the same positive recurrent set. The MLC movement *algorithm*, while no longer consistent with strict likelihood maximization outside the range $1/4 < x < 1/3$, is nevertheless a well defined movement rule and its performance is evaluated in the same fashion as that of the other algorithms in Figure 2.13. The hybrid strategy and the hybrid movement algorithm remain consistent over the range $0 \leq x < 1/3$. When $x > 1/3$ the MLC, ITC and mMLC strategies coincide with the ITC algorithm.

3.4 Relation to Ideal Observer Analysis

Our analysis of the present model is closely related to Ideal Observer Analysis [11]. Ideal observer analysis is a framework for quantitative assessment of the task-specific performance of sensory systems that finds wide use in psychophysics [12]. It complements information theory, with which it is sometimes confused; in the theoretical neuroscience literature the two approaches are often assumed to lead to the same conclusions, although there are important examples where they give different results [37]. In the results we show here, both the pursuit of food and the pursuit of information can be construed in terms of “ideal observers” with differing objectives. Our results concur with those of [11, 34, 37] and others in emphasizing the importance of task-specific objectives for characterizing “optimal” behavior of a system; information maximization in the absence of meaningful biological constraints is generally *not* sufficient as an explanatory principle for understanding biological behavior.

3.5 Future Directions

Each of the additional model elements mentioned in Section 3.2 can be taken into account while preserving the Markovian structure of the models considered. However, some of them would complicate the model sufficiently to require extensive computational analysis rather than exact treatment along the lines performed here.¹⁴ Whether through analytic or computational means, there is great opportunity for expanding our understanding of the interaction of a creature with a probabilistic environment through embedded Markovian models. Some directions that are accessible to analysis like that presented here include:

- Incorporating metabolic thresholds *e.g.* for starvation (a lower threshold on net metabolic resources) or for reproduction (setting an upper threshold to be reached before the creature can reproduce). Provided the statistical interaction of the creature and its sensorium preserves its Markovian character, the addition of an absorbing boundary allows the system to be analyzed in terms of first passage time distributions, for which there are abundant classical tools available. Assuming the creature is aware of its own metabolic resources and thresholds and adjusts its strategy accordingly, the creature’s internal reserve would become an additional component of its internal state.
- Allowing for multiple food types and heterogeneous metabolic resource requirements. Suppose a creature required two nutrients (“A” and “B”), and survival or reproduction depended on avoiding or encountering some boundary in the joint space representing internal reserves of both A and B. Suppose in addition the creature had limited sensing capacity. For instance, it might only be able to produce a finite total quantity of cell surface receptors divided between A and B. Clearly if it were near the threshold for mortality-for-want-of-A, it might choose to reduce the information it took in about B in order to increase its sensitivity to A, despite a net loss in overall information. Within this framework one may study precisely questions about the tradeoffs between net information and “relevant” information, a long standing conundrum in the application of information theory to biological systems.

¹⁴For example the action-perception system studied by Klyubin *et al.* [17] enjoys a Markovian structure, yet required computational rather than direct analysis.

- Motile cells and multicellular organisms often navigate by *chemotaxis*, detecting and climbing gradients of signaling molecules carrying information about the location of food or conspecific organisms. Chemotaxis based on gradient sensing is an example of a greedy algorithm that locally maximizes the expected rate of gain in the signal. As mentioned in the Introduction, Shraiman and colleagues introduced an alternative search strategy, *infotaxis*, that locally maximizes the expected rate of information gain rather than gain in food or signal concentration in a two dimensional geometry [40], and proposed it as an explanation for the zigzagging structure of casting paths observed in the flights of moths pursuing sparsely distributed chemoattractants. The information maximization strategy we consider here is formally equivalent to infotaxis. While we do not focus on the low concentration limit considered in [40], it is worth noting that in our system infotaxis outperforms a greedy concentration maximization strategy (the Max-Likelihood strategy) *only in the case when the source mobility is large, i.e. the source movement is relatively unpredictable*. Shraiman and colleagues only considered the case of a stationary source rather than a source performing a random walk. It would be interesting to investigate the case of a randomly moving source of a diffusible signal in a two dimensional geometry.

Is there a globally optimal strategy for food gathering for the system considered here? We have defined a Markovian algorithm as any map from the set of PMFs representing the creature’s current information about the source, to the creature’s next move. Any such strategy is equivalent to a Markov process on the state space of creature and source locations, observations and PMFs. Such a process can be described entirely by its transition matrix. If one considered the transition matrix on information states as a collection of free parameters, one could seek the optimal strategy relative to any objective function (such as maximizing the expected food gained over the long term, or the expected survival duration by some criterion) by numerically or analytically varying the transition matrix. It would be of great interest to compare the resulting optimal movement algorithm with those defined by heuristic strategies such as the infomax, track-and-pounce, and max likelihood strategies. As the complications considered above are more fully taken into account it may become possible to generate testable hypotheses about environmental scenarios under which selection pressures might lead to the evolution of distinct strategies, observable in actual organisms, related to the strategies investigated theoretically here.

4 Acknowledgments

PJT acknowledges support of the National Science Foundation, grant DMS-0720142, and research support from the Oberlin College Libraries. The authors thank K. Loparo for invaluable advice and discussion and for providing critical comments on a draft of the paper. The authors also wish to thank M. Lewicki, E. Meckes, B. Olshausen and T. Schneider for helpful discussion.

A Summary of Notation and Abbreviations

Essential notation and abbreviations, in (roughly) order of appearance:

N	Number of locations in the ring world.
\mathbb{Z}_N	The ring world, $\{0, 1, \dots, N-1\}$, with addition mod N .
t	Time: $t \in \{0, 1, 2, \dots\}$.
S_t	Source location at time t .
C_t	Creature location at time t .
PMF	Probability Mass Function.
P, P_x	PMF describing movement of the source.
x	Source mobility parameter, $1/4 < x < 1/3$.
$\mathcal{P}(\mathbb{Z}_N)$	Set of PMFs on \mathbb{Z}_N .
γ	Food injection rate.
α	Food decay rate.
y	Food molecule transition rate between adjacent nodes (fast timescale).
λ_m	Mean number of food molecules m steps CW from source.
Λ	Vector containing mean numbers of food molecules.
Z_t	Creature's observation at time t : either distance (D_t) or number (M_t).
\mathbf{Z}	The space of individual observations, <i>e.g.</i> $M_t \in \mathbb{N}$; $D_t \in \{0, 1, \dots, \lfloor N/2 \rfloor\}$.
$\lfloor u \rfloor$	The greatest integer less than or equal to a number u .
\mathcal{H}_t	The creature's history (of locations and observations) as of time t .
\mathcal{H}	The collection of all (infinitely long) histories.
F_t^+	PMF of source location at time t , given observations through time t .
f^+	A particular distribution taken (at random) by some F^+ .
F_t^-	PMF of source location at time t , given observations through time $t-1$.
f^-	A particular distribution taken (at random) by some F^- .
F_0^-	Initial PMF of source location, before any observations. $F_0^- \equiv 1/N$.
\mathcal{J}^+	The set of PMFs for which the injector's location is known, or is known to be some definite distance $d > 0$ from the creature's location.
\mathcal{J}^-	The set of PMFs obtained from $F \in \mathcal{J}^+$ by convolving with P .
\mathcal{J}^{a-}	Particular information states ($a = 0, 1, 2, 3, I$ or $*$).
ITC	Information theory Creature: creature following infomax strategy.
MLC	Max-Likelihood Creature: creature following likelihood maximization strategy.
mMLC	Hybrid or modified Max-Likelihood Creature: creature following hybrid strategy.
π	A stationary distribution of a Markov process.

B Food Distribution

The distribution of food molecules in Model I is obtained as the steady state of a rapidly equilibrating discrete time stochastic process. The following food distribution process is assumed to converge to equilibrium much more quickly than the time scale for source or creature movements: (i) Molecules of food enter the world at the location of the food source *via* a Poisson process with mean γ . After injection they transition between adjacent nodes with probability y . Molecules leave the world with probability α uniformly in space and time. We assume both $0 < y, \alpha < 1$. It can be shown (see [2], Chapter 3) that the molecule counts at each location relative to the injector location are independent and Poisson distributed with means given by a vector Λ . We obtain this

vector as follows. Let Q_y and R_γ be respectively the transition matrix and constant source vector

$$Q_y = \begin{pmatrix} 1-2y & y & & y \\ y & 1-2y & & \\ & & \ddots & \\ y & & & 1-2y \end{pmatrix} \quad R_\gamma = \begin{pmatrix} \gamma \\ 0 \\ \vdots \\ 0 \end{pmatrix}.$$

Then Λ satisfies $\Lambda = (1 - \alpha)Q_y\Lambda + R_\gamma$. Equivalently,

$$((1 - \alpha)Q_y - I)\Lambda = -R_\gamma. \quad (\text{B.1})$$

Lemma B.1. *The matrix $(1 - \alpha)Q - I$ is non-singular and hence Equation B.1 has unique solution Λ .*

Proof. $\sum_m Q_{lm} = 1$ because every molecule that was just moved must have come from some location with a probability of 1. And $1 > 1 - \alpha > 0$ because $1 > \alpha > 0$. Therefore, $1 > (1 - \alpha)\sum_m Q_{lm} > 0$. Subtracting $(1 - \alpha)Q_{ll}$ from both sides of $1 > (1 - \alpha)\sum_m Q_{lm}$ and taking the absolute value results in $|(1 - \alpha)Q_{ll} - 1| > (1 - \alpha)|\sum_{m:l \neq m} Q_{lm}|$. With the triangle inequality $|(1 - \alpha)Q_{ll} - 1| > (1 - \alpha)\sum_{m:l \neq m} |Q_{lm}|$. Each diagonal entry of the matrix $((1 - \alpha)Q - I)$ is $(1 - \alpha)Q_{ll} - 1$ and the sum of all entries of any row excluding the diagonal is $(1 - \alpha)\sum_{m:l \neq m} |Q_{lm}|$. Therefore $(1 - \alpha)Q - I$ is a diagonally dominant matrix. By the Levy-Desplanques theorem [36], $(1 - \alpha)Q - I$ is non-singular. \square

Because $(1 - \alpha)Q_y - I$ is non-singular, it is invertible. Therefore,

$$\Lambda = -((1 - \alpha)Q_y - I)^{-1}R_\gamma. \quad (\text{B.2})$$

Considering Λ as a function of γ ,

$$\begin{aligned} \Lambda(\gamma) &= -((1 - \alpha)Q_y - I)^{-1}R_\gamma \\ &= \gamma [-((1 - \alpha)Q_y - I)^{-1}R_1] \\ &= \gamma \Lambda(1). \end{aligned}$$

Due to the symmetry of Q_y and molecule injection occuring only at the source's location, Λ is symmetric about the source location. Therefore the expected molecule counts m steps CW and CCW of the injector are the same.

C Information Theory

Shannon's *Mathematical Theory of Communication* quantifies the uncertainty of a discrete random variable X taking values x with probability $\mathbb{P}(X = x)$ in terms of Boltzmann's *entropy* function ([32], Section I.6)

$$H(X) := - \left(\sum_{\{x: \mathbb{P}(X=x) \neq 0\}} \mathbb{P}(X = x) \log_2(\mathbb{P}(X = x)) \right), \quad (\text{C.1})$$

measured here in bits. When one variable depends conditionally on another we speak of the *conditional entropy of X given Y* :

$$\begin{aligned} H(X|Y) &:= \mathbb{E}_Y(H(X|Y = y)) \\ &= \sum_y \mathbb{P}(Y = y) H(X|Y = y). \end{aligned} \tag{C.2}$$

Observing the variable Y may or may not decrease our uncertainty about X , but it cannot (on average) increase it. Hence we have the inequality $H(X|Y) \geq H(X)$. A quantity of central interest in information theory is the *mutual information* of two random variables. It quantifies the (average) gain in information of X obtained from an observation of Y . Equivalently, it measures the departure from statistical independence of the random variables X and Y one another. The mutual information is given by

$$H(X; Y) := H(X) - H(X|Y). \tag{C.3}$$

In Section 2 we apply this expression to measure the creature’s information about the source location. In that case we take $H(X)$ to be the distribution of the source unconditioned on any observations, namely the uniform distribution on the ring, with entropy of $\log_2(N)$ bits. We take the conditional entropy $H(X|Y)$ to be the creature’s uncertainty about the source’s next location after making an observation, *i.e.* the distribution F^- , defined in Appendix D.2.

D Histories and Internal States

D.1 Histories

The creature searching for the source knows only the history of its own locations (relative to its starting position) and its observations of the local food concentration. In Model I these observations take the form of molecule counts at the creature’s location; in Model II the creature observes the inferred distance to the food source. Let the random variable Z_t represent the observation at time t in either case. Before making the observation, the creature knows its location but not how many molecules are present, or how distant the source is.

Technically, each history is one of an infinite set of all possible histories, a typical element of which would be

$$h = \{C_0 = c_0, Z_0 = z_0, C_1 = c_1, Z_1 = z_1, \dots, C_t = c_t, Z_t = z_t, \dots\}$$

(recall, however that we set $C_0 \equiv 0$ with probability one). When we consider the information known to the creature up to a certain point we will use conditional probabilities, conditioning on the set of histories specified up to the given time.¹⁵ For the argument in the main text it suffices to consider histories as lists of observations-to-date, but for completeness we give here a more formal definition, which is necessary to state the definition of the random measures F_t^\pm precisely.

¹⁵Our usage parallels the language of a filtration, an increasing family of σ -algebras in the theory of continuous time stochastic processes. The associated measure theoretic machinery will not be needed for the arguments here and we will simplify the notation as much as possible.

Let \mathcal{H} be the set of all histories of the form h given above, and \mathbf{Z} the space of individual observations ($\mathbf{Z} = \mathbb{N}$ in Model I, and $\mathbf{Z} = \{0, 1, \dots, \lfloor N/2 \rfloor\}$ in Model II). We define projection operators π_t^c and π_t^z from \mathcal{H} to \mathbb{Z}_N and \mathbf{Z} , respectively, as

$$\begin{aligned}\pi_t^c : h \in \mathcal{H} &\longrightarrow c_t \\ \pi_t^z : h \in \mathcal{H} &\longrightarrow z_t.\end{aligned}$$

We define *the history at time t* as the (random) subset of all histories consistent with the observations through time t :

$$\mathcal{H}_t = \{h \in \mathcal{H} | \pi_0^c(h) = C_0, \pi_0^z(h) = Z_0, \dots, \pi_t^c(h) = C_t, \pi_t^z(h) = Z_t\}.$$

On each time step t the latest creature location and observation is added *e.g.* :

$$\begin{aligned}\mathcal{H}_0 &= \{h \in \mathcal{H} | \pi_0^c(h) = C_0, \pi_0^z(h) = Z_0\} \\ \mathcal{H}_1 &= \{h \in \mathcal{H} | \pi_0^c(h) = C_0, \pi_0^z(h) = Z_0, \pi_1^c(h) = C_1, \pi_1^z(h) = Z_1\} \\ &\vdots\end{aligned}$$

Consequently the histories form a system of nested subsets of \mathcal{H} :

$$\mathcal{H} \supset \mathcal{H}_0 \supset \mathcal{H}_1 \supset \dots \supset \mathcal{H}_{t-1} \supset \mathcal{H}_t \supset \mathcal{H}_{t+1} \supset \dots$$

and we can form successive histories by intersection

$$\mathcal{H}_t = \mathcal{H}_{t-1} \cap \{h \in \mathcal{H} | \pi_t^c(h) = C_t, \pi_t^z(h) = Z_t\}. \quad (\text{D.1})$$

D.2 Internal States, or the Random Measures F^\pm

Given a history corresponding to the random sequences $\{C_t\}$ and $\{Z_t\}$, we can define what an ideal observer would be able to infer about the location of the source given the information available to the creature in terms of two probability mass functions. One (F_t^+ , say) captures what a creature knows about the current location of the source upon making the observation Z_t at time t . The other (F_{t+1}^- , say) captures what a creature knows about the *next* location of the source, before the creature itself moves and makes another observation. Hence F_{t+1}^- is determined entirely by F_t^+ and the rule governing the movement of the source. Formally, we may write

$$F_t^-(l) = \mathbb{P}(S_t = l | \mathcal{H}_{t-1}) \quad (\text{D.2})$$

$$F_t^+(l) = \mathbb{P}(S_t = l | \mathcal{H}_t) \quad (\text{D.3})$$

for $0 \leq l \leq N-1$. Hence F_t^\pm are random vectors that are also measures, or elements of the probability simplex on the ring

$$\mathcal{P}(\mathbb{Z}_N) = \{f \in \mathbb{R}^N | f \geq 0, \sum_i f_i = 1\}.$$

E Bayesian Update of Probability Mass Functions

Here we derive Equation 2.2 giving the Bayesian rule that updates the probability mass function F_t^+ , which is the PMF of the source's location at time t given the creature's history. We start from the definition of F_t^+ given by Equation D.3. For the conditional probability $P(S_t = l | h \in \mathcal{H}_t)$ to be defined we restrict consideration to histories \mathcal{H}_t that have strictly positive possibility of occurring, *i.e.* we restrict \mathcal{H}_t such that $\mathbb{P}(h \in \mathcal{H}_t) > 0$. To consolidate notation, let \mathcal{S} denote the event $S_t = l$, \mathcal{C} denote the event that $C_t = c$, \mathcal{Z} denote the event that $Z_t = z$, and H_t denote the event $h \in \mathcal{H}_t$. In this notation, note that $H_t = \mathcal{Z} \cap \mathcal{C} \cap H_{t-1}$ (*cf.* Equation D.1). Applying Bayes law, $\mathbb{P}(A \cap B) = \mathbb{P}(A|B)\mathbb{P}(B)$, Equation D.3 for $F_t^+(s)$ can be expanded as follows:

$$\begin{aligned} F_t^+(l) &= \mathbb{P}(\mathcal{S}|H_t) \\ &= \frac{\mathbb{P}(\mathcal{S} \cap \mathcal{Z} \cap \mathcal{C} \cap H_{t-1})}{\mathbb{P}(\mathcal{Z} \cap \mathcal{C} \cap H_{t-1})} \end{aligned} \tag{E.1}$$

$$= \frac{\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap \mathcal{S} \cap H_{t-1}) \cdot \mathbb{P}(\mathcal{S}|\mathcal{C} \cap H_{t-1}) \cdot \mathbb{P}(\mathcal{C} \cap H_{t-1})}{\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap H_{t-1}) \cdot \mathbb{P}(\mathcal{C} \cap H_{t-1})} \tag{E.2}$$

$$= \frac{\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap \mathcal{S} \cap H_{t-1}) \cdot \mathbb{P}(\mathcal{S}|\mathcal{C} \cap H_{t-1})}{\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap H_{t-1})}. \tag{E.3}$$

The number of molecules at C_t is a Poisson random variable with mean $\lambda_{S_t-C_t}$ and the distance D_t is $S_t - C_t$, therefore both observations are conditionally independent of the history given $C_t = c$ and $S_t = l$. But because the creature's observation Z_t is conditionally independent of the history given $C_t = c$ and $S_t = l$, we have

$$\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap \mathcal{S} \cap H_{t-1}) = \mathbb{P}(\mathcal{Z}|\mathcal{C} \cap \mathcal{S}).$$

The injector's location is conditionally independent of the creature's location C_t given \mathcal{H}_{t-1} because $S_t = S_{t-1} * P$ and C_t is a mapping from \mathcal{H}_{t-1} to \mathbb{Z}_N ; this fact implies

$$\begin{aligned} \mathbb{P}(\mathcal{S}|\mathcal{C} \cap H_{t-1}) &= \mathbb{P}(\mathcal{S}|H_{t-1}) \\ &= F_t^-(l). \end{aligned}$$

The second equality follows from the definition of F_t^- in Appendix D.2. Applying Bayes law and both examples of conditional independence to the term $\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap H_{t-1})$ we see,

$$\begin{aligned} \mathbb{P}(\mathcal{Z}|\mathcal{C} \cap H_{t-1}) &= \sum_{m=0}^{N-1} \mathbb{P}(\mathcal{Z}|S_t = m \cap \mathcal{C} \cap H_{t-1}) \cdot \mathbb{P}(S_t = m|\mathcal{C} \cap H_{t-1}) \\ &= \sum_{m=0}^{N-1} \mathbb{P}(\mathcal{Z}|S_t = m \cap \mathcal{C}) \cdot \mathbb{P}(S_t = m|H_{t-1}) \\ &= \sum_{m=0}^{N-1} \mathbb{P}(\mathcal{Z}|S_t = m \cap \mathcal{C}) \cdot F_t^-(m). \end{aligned}$$

Therefore, $F_t^+(l)$ can be written as

$$\begin{aligned} F_t^+(l) &= \frac{\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap \mathcal{S} \cap H_{t-1}) \cdot \mathbb{P}(\mathcal{S}|\mathcal{C} \cap H_{t-1})}{\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap H_{t-1})} \\ &= \frac{\mathbb{P}(\mathcal{Z}|\mathcal{C} \cap \mathcal{S}) \cdot F_t^-(l)}{\sum_{m=0}^{N-1} \mathbb{P}(\mathcal{Z}|S_t = m \cap \mathcal{C}) \cdot F_t^-(m)}, \end{aligned}$$

which concludes the derivation of Equation 2.2.

F Strategies

F.1 Strategies Defined in Terms of Probability Mass Functions and Information States

The Max-Likelihood strategy is defined as the strategy under which the creature moves to the injector's *most likely next location*, given the creature's history. If more than one location is equally likely, the MLC algorithm chooses the closest. If two are equally likely and equally close, the algorithm chooses between them at random. Given the next source location PMF F_{t+1}^- , define two subsets of locations $\mathcal{E}'_M \subset \mathcal{E}_M \subset \mathbb{Z}_N$ as

$$\mathcal{E}_M = \arg \max_{l \in \mathbb{Z}_N} F_{t+1}^-(l) \quad (\text{F.1})$$

$$\mathcal{E}'_M = \arg \min_{l' \in \mathcal{E}'_M} |C_t - l'| \quad (\text{F.2})$$

and choose $C_{t+1}|F_{t+1}^-$ at random, uniformly, from \mathcal{E}'_M .

Under the Infomax strategy the creature moves to a location that will give the lowest expected uncertainty of the source's location, upon the next observation. When more than one location would satisfy this criterion, the creature selects among the most informative locations those that provide the highest expected food reward. If multiple locations satisfy *both* the information maximization and utility maximization criteria, the creature selects the closest such location. If two locations are equally informative, food-rich, and close, the creature selects between them at random. Given the source location PMF F_{t+1}^- , define sets of locations $\mathcal{E}''_I \subset \mathcal{E}'_I \subset \mathcal{E}_I \subset \mathbb{Z}_N$ as

$$\mathcal{E}_I = \arg \max_{l \in \mathbb{Z}_N} \mathbb{E}[-H(F_{t+2}^-)|C_{t+1} = l] \quad (\text{F.3})$$

$$\mathcal{E}'_I = \arg \max_{l' \in \mathcal{E}_I} \mathbb{E}(M_{t+1}|C_{t+1} = l') \quad (\text{F.4})$$

$$\mathcal{E}''_I = \arg \min_{l'' \in \mathcal{E}'_I} |C_t - l''| \quad (\text{F.5})$$

Here $H(f)$ is the entropy of the PMF $f \in \mathcal{P}(\mathbb{Z}_N)$. Choose $C_{t+1}|F_{t+1}^-$ at random, uniformly, from \mathcal{E}''_I .

The modified Maximum Likelihood strategy (*a.k.a.* the Hybrid or the Track-and-Pounce strategy) chooses between the two previous strategies, depending on the creature's information state. When $F_{t+1}^- \in \mathcal{J}^{0-}$, the creature selects C_{t+1} from \mathcal{E}'_M , otherwise the creature selects C_{t+1} from \mathcal{E}''_I .

F.2 The Infomax Strategy Minimizes Uncertainty

The mutual information between the source's location at time t and the creature's prior history is determined by Equation C.3. This is equivalent to applying the same equation to each informa-

tion state. From Lemma F.1, we have the first term of Equation C.3 and we then calculate the conditional entropy of each State.

Lemma F.1. *The (unconditioned) probability that the injector is at location s at time t is $1/N$. That is, $\mathbb{P}(S^t = s) = 1/N$.*

Proof. The transition matrix P generating the injector's next location is independent of its current location. In addition, the initial distribution of the injector is assumed to be uniform on the ring, i.e. $\forall s \mathbb{P}(S_0 = s) = 1/N$. Consider the injector starting at some location, l , moving along some path according to P and ending up at location m . It is equally probable that the injector could have started at $l + a$ because

$$\mathbb{P}(S_0 = l) = \mathbb{P}(S_0 = l + a) = 1/N,$$

moved along the same (equally probable) path, and ended up at location $m + a$. This argument holds for all starting locations and paths so $\forall a, \forall t$

$$\mathbb{P}(S^t = s) = \mathbb{P}(S^t = s + a).$$

Therefore, $\forall t, \forall s, \mathbb{P}(S^t = s) = 1/N$. □

The infomax strategy greedily optimizes for expected information and then expected food. For the Distance-Certain Model we propose a strategy in Figure 2.7 which we now show is an infomax strategy. Furthermore, we demonstrate that the infomax strategy is long-term optimal for information. First we demonstrate that a creature in \mathcal{J}^{0-} has more information about the injector's location than a creature in any other state and then point out that the strategy proposed in Figure 2.7 always transitions to \mathcal{J}^{0-} . Finally, we show that the proposed strategy selects the location with highest expected food that is guaranteed to transition to \mathcal{J}^{0-} .

Lemma F.2. *A creature in the information state \mathcal{J}^{0-} has optimal information about the injector's location.*

Proof. Recall that \mathcal{J}^+ is the set of information states in which the distance to the source is known with certainty. We partition \mathcal{J}^+ into the following subsets: \mathcal{J}^{0+} where the injector's location is known, \mathcal{I}^{1+} where the injector is known to be one step away from some location but not in \mathcal{J}^{0+} , and \mathcal{I}^{2+} where the injector is known to be a distance d away from some location ($d \geq 2$) but not in \mathcal{J}^{0+} . This partition of \mathcal{J}^+ leads to a natural partition of \mathcal{J}^- where \mathcal{J}^{0-} , \mathcal{I}^{1-} , and \mathcal{I}^{2-} are generated from \mathcal{J}^{0+} , \mathcal{I}^{1+} , and \mathcal{I}^{2+} respectively by application of the convolution representing the injector's movement algorithm P_x . Note that $\mathcal{J}^{1-} \subset \mathcal{I}^{1-}$ because \mathcal{I}^{1-} also includes situations where the source is not equally likely to be CW as CCW.

Each f^0 in \mathcal{J}^{0-} has the form

$$f^0(l) = \begin{cases} x & l = e - 1 \\ 1 - 2x & l = e \\ x & l = e + 1 \\ 0 & \text{otherwise} \end{cases} \quad (\text{F.6})$$

with $e \in \{0, 1, \dots, N-1\}$.

Each f^1 in I^{1-} has the form

$$f^1(l) = \begin{cases} px & l = e - 2 \\ p(1 - 2x) & l = e - 1 \\ x & l = e \\ q(1 - 2x) & l = e + 1 \\ qx & l = e + 2 \\ 0 & \text{otherwise} \end{cases} \quad (\text{F.7})$$

with $e \in \{0, 1, \dots, N-1\}$, $p, q > 0$, and $p + q = 1$. Here p is the probability that the injector was 1 step CCW before moving.

Each f^2 in I^{2-} has the form

$$f^2(l) = \begin{cases} px & l = e - d \pm 1 \\ p(1 - 2x) & l = e - d \\ qx & l = e + d \pm 1 \\ q(1 - 2x) & l = e + d \\ 0 & \text{otherwise} \end{cases} \quad (\text{F.8})$$

with $e \in \{0, 1, \dots, N-1\}$, $d \in \{2, 3, \dots, N-2\}$, $p, q > 0$, and $p + q = 1$. Here p is the probability that the injector was $1 < d < N-1$ steps CCW before moving.

If a creature has a history that maps F^- into \mathcal{J}^{0-} , then the mutual information between the injector's location and the creature's prior history is given by Equation C.3 as

$$\begin{aligned} H_0(S_t; h \in \mathcal{H}_{t-1} \cap h \mapsto \mathcal{J}^{0-}) &= H(S^t) - H(S^t | h \in \mathcal{H}_{t-1} \cap h \mapsto \mathcal{J}^{0-}) \\ &= H(S^t) - \mathbb{E}_{f^0 \in \mathcal{J}^{0-}} [H(F_t^- = f^0)] \\ &= \log_2(N) + 2x \log_2(x) + (1 - 2x) \log_2(1 - 2x). \end{aligned}$$

For convenience, let us define the function $\phi : [0, 1] \mapsto [0, 1/2]$ as

$$\phi(x) = \begin{cases} 0, & x = 0 \\ x \log_2(x), & 0 < x \leq 1. \end{cases} \quad (\text{F.9})$$

If a creature has a history that maps F^- into I^{1-} , then the mutual information between the injector's location and the creature's prior history (given by Equation C.3) depends on the CCW probability p , $0 < p < 1$. (If $p = 0$ or $p = 1$ then we are back in information state \mathcal{J}^{0+} , as the location of the source is known with certainty). As a function of p , the mutual information is

$$\begin{aligned} H_1(S_t; h \in \mathcal{H}_{t-1} \cap h \mapsto I^{1-}) &= H(S^t) - H(S^t | h \in \mathcal{H}_{t-1} \cap h \mapsto I^{1-}) \\ &= \log_2(N) + 2x \log_2(x) + (1 - 2x) \log_2(1 - 2x) + \phi(p)(1 - x) \\ &= H_0(S_t; h \in \mathcal{H}_{t-1} \cap h \mapsto \mathcal{J}^{0-}) + \phi(p)(1 - x). \end{aligned}$$

Similarly, if a creature has a history that maps F^- into I^{2-} , then the mutual information between the injector's location and the creature's prior history is

$$\begin{aligned} H_2(S_t; h \in \mathcal{H}_{t-1} \cap h \mapsto I^{2-}) &= H(S^t) - H(S^t | h \in \mathcal{H}_{t-1} \cap h \mapsto I^{2-}) \\ &= \log_2(N) + 2x \log_2(x) + (1 - 2x) \log_2(1 - 2x) + \phi(p) \\ &= H_0(S_t; h \in \mathcal{H}_{t-1} \cap h \mapsto \mathcal{J}^{0-}) + \phi(p). \end{aligned}$$

As we have restricted the source mobility parameter x to the range $1/4 \leq x \leq 1/3$, we have the strict inequality that for each p , $H_0 > \max(H_1, H_2)$. That is, the mutual information (the reduction in entropy of S_t upon conditioning on the creature's history \mathcal{H}_{t-1}) is greatest for a creature in information state \mathcal{J}^{0-} . As \mathcal{J}^{0-} , \mathcal{I}^{1-} , and \mathcal{I}^{2-} partition all possible histories in the Distance Certain Model (together with the initial state f_u which represents a state of having no information about the source), a creature in the information state \mathcal{J}^{0-} has the maximal possible information about the injector's location. \square

Corollary F.3. *The Infomax Strategy has (long-term) optimal information about the injector's location.*

Proof. From Lemma F.2, the information state \mathcal{J}^{0-} minimizes the creature's uncertainty about the source's location, *i.e.* it minimizes the entropy of $f \in \mathcal{J}^-$. The infomax strategy illustrated in Figure 2.7 transitions to \mathcal{J}^{0-} with probability one (see Figure 2.8). Therefore, with probability one, the infomax strategy has optimal information about the injector's location. \square

The infomax strategy must choose between locations that have optimal expected information on the next turn. Lemma F.2 shows that a location that is guaranteed to transition to \mathcal{J}^{0-} has the highest possible expected information. The infomax strategy proposed in Figure 2.7 chooses the location that is guaranteed to transition to \mathcal{J}^{0-} and has the highest expected food of all locations that transition to \mathcal{J}^{0-} . Therefore Figure 2.7 depicts an infomax strategy.

F.2.1 Mutual Information for Sequence Entropies

In addition to calculating the mutual information as the expected reduction in entropy between $f_u \equiv 1/N$ and F^- , one may consider the differences in entropies between *ensembles of random sequences* of source locations on the one hand, and creature locations and internal states on the other. One may compare the entropy of the finite sequence (S_0, S_1, \dots, S_T) unconditioned on any observations with the mean entropy after conditioning on the creature's trajectory (C_0, C_1, \dots, C_T) and its sequence of internal states $(F_0^+, F_1^+, \dots, F_T^+)$. Allowing for a slight abuse of notation, let the (negative) entropy function ϕ defined above be written with three arguments in the case of three possible outcomes: we will write $\phi(a, b, c)$ for $a \log_2 a + b \log_2 b + c \log_2 c$, for example, or $\phi(x, 1 - 2x, x)$ for $2x \log_2(x) + (1 - 2x) \log_2(1 - 2x)$. It is easy to show, using chains of conditional probabilities, that the entropy of the finite sequence of source locations (S_0, S_1, \dots, S_T) is given by

$$H[\{S_t\}_{t=0}^T] = \log_2(N) - T\phi(x, 1 - 2x, x) \quad (\text{F.10})$$

where the $\log_2(N)$ term reflects the entropy due to the initial (uniform) distribution and the $\phi(x, 1 - 2x, x)$ terms reflect the contribution to the entropy of the finite sequence on each successive time step.¹⁶ We may define the *sequence entropy* to be the mean growth rate of the entropy of the

¹⁶We note that this construction works precisely because the source alone – like the entire system – is Markov process.

sequence over time,¹⁷ *i.e.*

$$H_{seq} = \lim_{T \rightarrow \infty} \left(\frac{1}{T} H[\{S_t\}_{t=0}^T] \right) = \lim_{T \rightarrow \infty} \left(\frac{1}{T} [\log_2(N) - T\phi(x, 1-2x, x)] \right) = -\phi(x, 1-2x, x). \quad (\text{F.11})$$

As expected for any Markov process that tends towards an equilibrium distribution, the entropy of the sequence T steps after the initial condition asymptotically grows linearly in time. One may carry out a similar calculation for the entire system restricted to the set of one, two or three absorbing states (see *e.g.* Appendix F.3). It suffices to point out, however, that if the creature adopts the infomax strategy, hence returns on each step to the “most informed” of the information states \mathcal{J}^{0+} , then the *sequence of source locations conditioned on the creature’s history* has asymptotic entropy of zero. The creature that remains in the information state \mathcal{J}^{0+} can retrospectively reconstruct the source trajectory perfectly. The other creatures cannot, and hence the infomax strategy is information-optimal both from the point of view of the next-step entropies and of the entropy of the entire sequence of source trajectories and creature histories.

F.3 Transition Matrices

The state space of any finite Markov chain may be uniquely decomposed into a set of transient states and one or more irreducible sets of recurrent states. A state is *recurrent* if the probability of eventual return to that state is unity, otherwise the state is *transient*. A set of states is *irreducible* if all the states in the set intercommunicate, *i.e.* there is a positive probability of moving between any pair over a finite number of iterations. For example, if M is a transition matrix defined on a set of states of the Markov chain then the set is irreducible if every entry of M^k is positive for some k . An irreducible Markov chain on a finite state space will have a unique stationary distribution, π [14]. We will show that each of the strategies considered reduces to an irreducible Markov chain on a finite set of recurrent states. We call this subset of the original Markov chain the *absorbing set*. We will calculate the stationary distribution π in each case.

The transition matrix between information states for the Infomax strategy is trivial. The system enters the state \mathcal{J}^{0-} by the second iteration and remains in state \mathcal{J}^{0-} from then on, with probability 1. In Table 2.1, the matrix M_{ITC} gives the transitions restricted to the five relevant information states. State \mathcal{J}^{0-} corresponds to the second column, and is absorbing after two iterations for all initial conditions. Equivalently, the matrix M_{ITC}^2 consists of a single nonzero column containing all unit entries.

The Hybrid strategy has an absorbing set consisting of the states \mathcal{J}^{0-} and \mathcal{J}^{1-} . The transition matrix for the strategy restricted to these states is

$$M = \begin{bmatrix} 1-2x & 2x \\ 1 & 0 \end{bmatrix}. \quad (\text{F.12})$$

This set of states is irreducible and recurrent since M^2 has all positive entries in this case. The

¹⁷We note that Schreiber introduced a related notion, the *transfer entropy*, to quantify the transfer of information from one sequence of random variables to another [31].

stationary distribution is

$$\pi_{mMLC} = \begin{bmatrix} \frac{1}{1+2x} \\ \frac{2x}{1+2x} \end{bmatrix}. \quad (\text{F.13})$$

The Max-Likelihood strategy has an absorbing set consisting of the states \mathcal{J}^{0-} , \mathcal{J}^{1-} and \mathcal{J}^{2-} . The transition matrix restricted to these states is

$$M = \begin{bmatrix} 1-2x & 2x & 0 \\ x & 1-2x & x \\ 1-x & x & 0 \end{bmatrix}. \quad (\text{F.14})$$

In this case as well these three sets form an irreducible recurrent chain; the matrix M^2 again has all positive entries, for the range of x considered ($1/4 < x < 1/3$). The stationary distribution is

$$\pi_{MLC} = \begin{bmatrix} \frac{2-x}{4+x} \\ \frac{2}{4+x} \\ \frac{2x}{4+x} \end{bmatrix}. \quad (\text{F.15})$$

G The Expected Food and Information

The mean food intake is the expected value (under the asymptotic probability distribution for the Markov process on the absorbing set) of the molecule count. Since the system is ergodic, this expected value is equal (“almost surely”) to the long time average of the molecule count. That is

$$\lim_{T \rightarrow \infty} \frac{1}{T} \sum_{t=1}^T M_t = \mathbb{E}[M_t] \quad (\text{almost surely}).$$

For a given strategy, the expected food is then the dot product of the stationary distribution π and the expected food for each state. Similarly, the expected information is the dot product of the stationary distribution π and the information in each state.

Let $\phi(x)$ denote the (negative) entropy function $\phi(x) = 2x \log_2(x) + (1-2x) \log_2(1-2x)$ as in Appendix F.2.1.

Table G shows the entropies of the probability mass functions F^- representing the uncertainty of the injector’s next location in each information state. If $H_k(x)$ denotes the entropy of the PMF for state \mathcal{J}^{k-} (as a function of the injector’s mobility, x) then we have $H_0(x) < H_1(x) < H_2(x) = H_k(x), k \geq 2$.

The long-run average entropy of the information states for each strategy is obtained from Table G and the equilibrium distributions in Appendix F.3. Let \bar{H}_{ITC} , \bar{H}_{mMLC} and \bar{H}_{MLC} be the mean

Information State	Entropy of PMF F^-
\mathcal{J}^{0-}	$H_0(x) = -\phi(x)$
\mathcal{J}^{1-}	$H_1(x) = -\phi(x) + 1 - x$
\mathcal{J}^{2-}	$H_2(x) = -\phi(x) + 1$

Table G.1: Entropy of probability mass functions (PMFs) in different information states. The PMF representing the next location of the source when the creature is in information state $k \in \{0, 1, 2\}$, is given as a function of the source mobility x by $H_k(x)$. We restrict the source mobility to $1/4 < x < 1/3$. For $k > 2$ we have $H_k(x) = H_2(x)$. Clearly $H_0(x) < H_1(x) < H_2(x)$. Notation: for $0 < x < 1/2$, $\phi(x)$ is here defined as $\phi(x) = 2x \log_2(x) + (1 - 2x) \log_2(1 - 2x)$.

entropies of the internal state F^- for the infomax, hybrid and max-likelihood creatures, respectively. From a straightforward calculation we obtain:

$$\begin{aligned}
\bar{H}_{ITC} &= -\phi(x) \\
\bar{H}_{mMLC} &= -\phi(x) + 2x(1 - x)/(1 + 2x) \\
\bar{H}_{MLC} &= -\phi(x) + 2/(4 + x),
\end{aligned}$$

from which it follows (in the range $1/4 \leq x \leq 1/3$) that the infomax strategy gives the source location estimates with the lowest mean entropy:

$$\bar{H}_{ITC} < \bar{H}_{mMLC} < \bar{H}_{MLC}.$$

H Joint (Distance, Information State) Transitions

To understand the dependence of expected food gained on source mobility, we derive the *joint* transition probabilities between states representing both the distance D_t between the creature and the source at each time step and the creature's information state \mathcal{J}_t^- , for the information states that form the positive recurrent states for each strategy. Let the notation (n, k) denote the creature at time t being a distance $D_t = n$ from the food source and in information state \mathcal{J}_t^{k-} . The conditional probability $\mathbb{P}_S(D_t = n | \mathcal{J}_t^{k-})$ that a creature following strategy S ends up a distance n from the

source, given that the creature is in information state \mathcal{J}^{k-} , is as follows:

$$\mathbb{P}_{ITC}(D_t = n | \mathcal{J}_t^{k-}) = \begin{array}{c|c} (n, k) & (\cdot, 0) \\ \hline (0, \cdot) & x \\ (1, \cdot) & 1 - 2x \\ (2, \cdot) & x \end{array} \quad (\text{H.1})$$

$$\mathbb{P}_{mMLC}(D_t = n | \mathcal{J}_t^{k-}) = \begin{array}{c|cc} (n, k) & (\cdot, 0) & (\cdot, 1) \\ \hline (0, \cdot) & 1 - 2x & x/2 \\ (1, \cdot) & 2x & (1 - 2x)/2 \\ (2, \cdot) & 0 & x \\ (3, \cdot) & 0 & (1 - 2x)/2 \\ (4, \cdot) & 0 & x/2 \end{array} \quad (\text{H.2})$$

$$\mathbb{P}_{MLC}(D_t = n | \mathcal{J}_t^{k-}) = \begin{array}{c|ccc} (n, k) & (\cdot, 0) & (\cdot, 1) & (\cdot, 2) \\ \hline (0, \cdot) & 1 - 2x & x & (1 - 2x)/2 \\ (1, \cdot) & 2x & 1 - 2x & x \\ (2, \cdot) & 0 & x & 0 \\ (3, \cdot) & 0 & 0 & x/2 \\ (4, \cdot) & 0 & 0 & (1 - 2x)/2 \\ (5, \cdot) & 0 & 0 & x/2 \end{array}. \quad (\text{H.3})$$

To find the probability of being distance n from the injector we apply Bayes law, obtaining $\mathbb{P}_S(D = n) = \sum_k \mathbb{P}_S(D = n | \mathcal{J}_t^{k-}) \mathbb{P}_S(F^- \in \mathcal{J}^{k-})$. We sum the index k over the positive recurrent information states for strategy S . The probability of being in each information state for each creature is calculated in Appendix F.3, and the resulting probabilities for each combination of mobility and strategy are given in Table 2.3.

With rows correspond to states at time t and columns correspond to states at time $t + 1$, the joint

(distance, information) transition matrices on the positive recurrent states are:

$$T_{ITC} = \begin{array}{c|ccc} (n, k) & (0, 0) & (1, 0) & (2, 0) \\ \hline (0, 0) & x & 1 - 2x & x \\ (1, 0) & x & 1 - 2x & x \\ (2, 0) & x & 1 - 2x & x \end{array} \quad (\text{H.4})$$

$$T_{mMLC} = \begin{array}{c|ccccccc} (n, k) & (0, 0) & (1, 0) & (0, 1) & (1, 1) & (1, 2) & (1, 3) & (1, 4) \\ \hline (0, 0) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 \\ (1, 0) & 0 & 0 & \frac{x}{2} & \frac{1}{2}(1 - 2x) & x & \frac{1}{2}(1 - 2x) & \frac{x}{2} \\ (0, 1) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 \\ (1, 1) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 \\ (2, 1) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 \\ (3, 1) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 \\ (4, 1) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 \end{array} \quad (\text{H.5})$$

$$T_{MLC} = \begin{array}{c|cccccccccc} (n, k) & (0, 0) & (1, 0) & (0, 1) & (1, 1) & (2, 1) & (0, 2) & (1, 2) & (3, 2) & (4, 2) & (5, 2) \\ \hline (0, 0) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (1, 0) & 0 & 0 & x & 1 - 2x & x & 0 & 0 & 0 & 0 & 0 \\ (0, 1) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (1, 1) & 0 & 0 & x & 1 - 2x & x & 0 & 0 & 0 & 0 & 0 \\ (2, 1) & 0 & 0 & 0 & 0 & 0 & \frac{1-2x}{2} & x & \frac{x}{2} & \frac{1-2x}{2} & \frac{x}{2} \\ (0, 2) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (1, 2) & 0 & 0 & x & 1 - 2x & x & 0 & 0 & 0 & 0 & 0 \\ (3, 2) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (4, 2) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ (5, 2) & 1 - 2x & 2x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{array}. \quad (\text{H.6})$$

From these transition matrices we obtain the stationary probability vectors for the joint (information state, distance) distributions, from which one immediately obtains the fraction of time spent at each distance relative to the food source. In the case of the MLC each distance $0 \leq D \leq 5$ appears in exactly one joint state, so the matrix T_{MLC} may be viewed as defining a Markov process on the six recurring distances zero through five. In the case of the hybrid creature the distance zero occurs both in conjunction with the information state 0 and the information state 1, which prevents representation of this strategy's behavior as a Markov process solely in terms of distances. The steady state distributions for the *distances* for the hybrid and MLC strategies are given algebraically in Table 2.3, and graphically in Figure 2.12.

I Expected Distance Traveled

We illustrate the calculation of expected distance traveled and movement probability under the infomax algorithm. The other cases are similar, requiring only additional bookkeeping. Consider the transitions for the joint representation of information state and distance $|C_t - S_t|$ between creature and source for the ITC given by Equation H.4. Once this creature enters information state \mathcal{J}^{0-} it remains in that state, with only its position relative to the injector varying. The

creature following the IT algorithm moves on each time step to the location either one to the right or one to the left of the injector's last know position, because these locations minimize the expected uncertainty on the following time step, and maximize the expected food gained amongst locations with equal expected uncertainty. In addition, given two locations jointly optimal for information (first) and food (second) the ITC choses the closer of the two destinations. There are three cases to consider.

1. With probability $1 - 2x$ the source remains in place, or $S_{t+1} = S_t$. In this case the creature remains in place as well, or $C_{t+1} = C_t$, and the distance moved is $\Delta_t = |C_{t+1} - C_t| = 0$.
2. With probability x the source moves clockwise, or $S_{t+1} = S_t + 1$.
 - (a) With probability $x/2$, $C_t = S_t - 1$. In this case, the creature moves one step CW, and $C_{t+1} = C_t + 1$.
 - (b) With probability $x/2$, $C_t = S_t + 1$. In this case the injector moves (by chance) to colocalize with the creature. The creature then moves either CW or CCW with equal likelihood, in order to remain in state \mathcal{J}^{0-} . Therefore
 - i. With probability $x/4$, $C_{t+1} = C_t + 1$.
 - ii. With probability $x/4$, $C_{t+1} = C_t - 1$.
3. With probability x the source moves counterclockwise, or $S_{t+1} = S_t - 1$.
 - (a) With probability $x/2$, $C_t = S_t + 1$. In this case, the creature moves one step CCW, and $C_{t+1} = C_t - 1$.
 - (b) With probability $x/2$, $C_t = S_t - 1$. The creature then moves either CW or CCW with equal likelihood, so:
 - i. With probability $x/4$, $C_{t+1} = C_t + 1$.
 - ii. With probability $x/4$, $C_{t+1} = C_t - 1$.

In sum, the creature remains in information state \mathcal{J}^{0-} with probability 1 , and moves a distance $\Delta_t = 1$ with probability $2x$, or a distance $\Delta_t = 0$ with probability $1 - 2x$. Compare the middle column of Table 2.5. The calculations for the MLC and mMLC strategies follow in a similar fashion from the transition probabilities given in Equations H.5-H.6.

References

- [1] C. Adami. *Introduction to artificial life*. Springer-Verlag, 1998.
- [2] Edward K. Agarwala. Food for thought: When information optimization fails to optimize utility. Master's thesis, Case Western Reserve University, Cleveland, Ohio, June 2009.
- [3] W. Rosh Ashby. *An Introduction to Cybernetics*. Chapman & Hall LTD, 1956.
- [4] F. Attneave. Some informational aspects of visual perception. *Psychol. Rev.*, 61:183–93, 1954.
- [5] HB Barlow. *Sensory Communications*, chapter Possible principles underlying the transformation of sensory messages. MIT Press, 1961.

- [6] A J Bell. Levels and loops: the future of artificial intelligence and neuroscience. *Philos Trans R Soc Lond B Biol Sci*, 354(1392):2013–20, Dec 1999.
- [7] H J Chiel and R D Beer. The brain has a body: adaptive behavior emerges from interactions of nervous system, body and environment. *Trends Neurosci*, 20(12):553–7, Dec 1997.
- [8] X. Feng and K.A. Loparo. Active probing for information in control systems with quantized state measurements: a minimum entropy approach. *Automatic Control, IEEE Transactions on*, 42(2):216–238, feb. 1997.
- [9] Xiangbo Feng, K.A. Loparo, and Yuguang Fang. Optimal state estimation for stochastic systems: an information theoretic approach. *Automatic Control, IEEE Transactions on*, 42(6):771–785, jun. 1997.
- [10] Joaquín M Fuster. Upper processing stages of the perception-action cycle. *Trends Cogn Sci*, 8(4):143–5, Apr 2004.
- [11] Wilson S. Geisler. *The Visual Neurosciences*, chapter 52: Ideal Observer Analysis. MIT Press, 2004.
- [12] David M. Green and John A. Swets. *Signal detection theory and psychophysics*. Wiley, New York, 1966.
- [13] Barry E. Griffiths and Kenneth A. Loparo. Optimal control of jump-linear gaussian systems. *International Journal of Control*, 42:791–819, 1985.
- [14] Geoffrey Grimmett and David Stirzaker. *Probability and Random Processes*. Oxford University Press, third edition, 2005.
- [15] Devin L Jindrich and Robert J Full. Dynamic stabilization of rapid hexapedal locomotion. *J Exp Biol*, 205(Pt 18):2803–23, Sep 2002.
- [16] Michael I. Jordan. *Learning in Graphical Models*. MIT Press, 1998.
- [17] Alexander S Klyubin, Daniel Polani, and Chrystopher L Nehaniv. Representations of space and time in the maximization of information flow in the perception-action loop. *Neural Comput*, 19(9):2387–432, Sep 2007.
- [18] Donald L. Kramer and Robert L. McLaughlin. The behavioral ecology of intermittent locomotion. *American Zoologist*, 41(2):137–153, 2001.
- [19] R Kukillaya, J Proctor, and P Holmes. Neuromechanical models for insect locomotion: Stability, maneuverability, and proprioceptive feedback. *Chaos*, 19(2):026107, Jun 2009.
- [20] Ralph Linsker. Self-organization in a perceptual network. *Computer*, 21:105–117, 1988.
- [21] Sergey N Markin, Alexander N Klishko, Natalia A Shevtsova, Michel A Lemay, Boris I Pri-lutsky, and Ilya A Rybak. Afferent control of locomotor cpg: insights from a simple neuromechanical model. *Ann N Y Acad Sci*, 1198:21–34, Jun 2010.
- [22] J M McNamara and A I Houston. State-dependent life histories. *Nature*, 380(6571):215–21, Mar 1996.

- [23] Jiri Najemnik and Wilson S Geisler. Optimal eye movement strategies in visual search. *Nature*, 434(7031):387–91, Mar 2005.
- [24] W. John O’Brien, Barbara I. Evans, and Howard I. Browman. Flexible search tactics and efficient foraging in saltatory searching animals. *Oecologia*, 80:100–110, 1989.
- [25] Charles Ofria, Christoph Adami, and Travis C Collier. Selective pressures on genomes in molecular evolution. *J Theor Biol*, 222(4):477–83, Jun 2003.
- [26] Daniel Polani. Information: currency of life? *Human Frontier Science Program Journal*, 8 Sept 2009.
- [27] Andy M Reynolds and Mark A Frye. Free-flight odor tracking in *Drosophila* is consistent with an optimal intermittent scale-free search. *PLoS One*, 2(4):e354, 2007.
- [28] T. D. Schneider. Evolution of biological information. *Nucleic Acids Res.*, 28(14):2794–2799, 2000.
- [29] T. D. Schneider, G. D. Stormo, L. Gold, and A. Ehrenfeucht. Information content of binding sites on nucleotide sequences. *J. Mol. Biol.*, 188:415–431, 1986.
- [30] Thomas D Schneider. 70% efficiency of bistate molecular machines explained by information theory, high dimensional geometry and evolutionary convergence. *Nucleic Acids Res*, Jun 2010.
- [31] Schreiber. Measuring information transfer. *Phys Rev Lett*, 85(2):461–4, Jul 2000.
- [32] C.E. Shannon. A mathematical theory of communication. *The Bell System Technical Journal*, 27:379–423,623–656, July, October 1948.
- [33] M L Shik and G N Orlovsky. Neurophysiology of locomotor automatism. *Physiol Rev*, 56(3):465–501, Jul 1976.
- [34] E P Simoncelli and B A Olshausen. Natural image statistics and neural representation. *Annu Rev Neurosci*, 24:1193–216, 2001.
- [35] E.C. Smith and M.S. Lewicki. Efficient Auditory Coding. *Nature*, 439(7079):978–982, 2006.
- [36] O. Taussky. A recurring theorem on determinants. *American Mathematical Monthly*, 56(4):672–676, 1949.
- [37] Eric E Thomson and William B Kristan. Quantifying stimulus discriminability: a comparison of information theory and ideal observer analysis. *Neural Comput*, 17(4):741–778, Apr 2005.
- [38] Touchette and Lloyd. Information-theoretic limits of control. *Phys Rev Lett*, 84(6):1156–9, Feb 2000.
- [39] J.H. van Hateren and A. van der Schaaf. Independent component filters of natural images compared with simple cells in primary visual cortex. *Proceedings: Biological Sciences*, 265(1394):359–366, 1998.
- [40] M. Vergassola, E. Villermanx, and B.I. Shraiman. ‘Infotaxis’ as a strategy for searching without gradients. *Nature*, 445(7126):406–409, 2007.
- [41] Martin J. Wainwright. *Stochastic processes on graphs with cycles: geometric and variational approaches*. PhD thesis, Massachusetts Institute of Technology, 2002.

- [42] M.J. Wainwright and M.I. Jordan. Log-determinant relaxation for approximate inference in discrete markov random fields. *Signal Processing, IEEE Transactions on*, 54(6):2099 – 2109, june 2006.